arXiv:1611.06173v1 [math.ST] 18 Nov 2016

EMPIRICAL RISK MINIMIZATION AND COMPLEXITY OF DYNAMICAL MODELS

KEVIN MCGOFF, ANDREW B. NOBEL

Abstract. A dynamical model consists of a continuous self-map $T : \mathcal{X} \to \mathcal{X}$ of a compact state space $\mathcal{X}$ and a continuous observation function $f : \mathcal{X} \to \mathbb{R}$. Dynamical models give rise to real-valued sequences when f is applied to repeated iterates of T beginning at some initial state. This paper considers the fitting of a parametrized family of dynamical models to an observed real-valued stochastic process using empirical risk minimization. More precisely, at each finite time, one selects a model yielding values that minimize the average per-symbol loss with the observed process. The limiting behavior of the minimum risk parameters is studied in a general setting where the observed process need not be generated by a dynamical model in the family and minimal conditions are placed on the loss function.

We establish a general convergence theorem for minimum risk estimators and ergodic observations, showing that the limiting parameter set is characterized by the projection of the observed process onto the set of processes associated with the dynamical family, where the projection is taken with respect to a joining-based distortion between stationary processes that is a generalization of the $\overline{d}$-distance. We then turn our attention to observations generated from an ergodic process plus an i.i.d. noise process and study conditions under which empirical risk minimization can effectively separate the signal from the noise. The key, necessary condition in the latter results is that the family of dynamical models has limited complexity, which is quantified through a notion of entropy for families of infinite sequences. Close connections between entropy and limiting average mean widths for stationary processes are established.

1. Introduction

Empirical risk minimization is a common approach to model fitting and estimation in a variety of parametric and non-parametric problems. In this paper we investigate the use of empirical risk minimization to fit a family of dynamical models to an observed stochastic process. Formally, a dynamical model consists of a continuous transformation $T : \mathcal{X} \to \mathcal{X}$ on a compact metric space $\mathcal{X}$, and a continuous observation function $f : \mathcal{X} \to \mathbb{R}$. Let T^k denote the k-fold composition of T with itself, and let T^0 be the identity map on $\mathcal{X}$. From each initial state $x \in \mathcal{X}$ the dynamical model (T, f) yields a real-valued sequence

$$\{f(T^k x) \,=\, f \circ T^k(x) : k \geq 0\} \,\subseteq\, \mathbb{R},$$

obtained by applying the observation function f to a deterministic sequence of states generated by repeated iteration of the transformation T. In general, f need not be injective, so one cannot necessarily recover the underlying states from the sequence $\{f(T^k x)\}_k$.

In what follows we consider an indexed family $\mathcal{D} = \{(T_\theta, f_\theta) : \theta \in \Theta\}$ of dynamical models defined on a common compact metric space $\mathcal{X}$ satisfying the following conditions:

(D1) the index set Θ is a compact metric space;
(D2) the map $(\theta, x) \mapsto T_\theta(x)$ from $\Theta \times \mathcal{X}$ to $\mathcal{X}$ is continuous;
(D3) the map $(\theta, x) \mapsto f_\theta(x)$ from $\Theta \times \mathcal{X}$ to $\mathbb{R}$ is continuous.

Condition (D2) ensures that each transformation T_θ is continuous and that the action of T_θ is continuous in θ. Condition (D3) ensures that each observation function f_θ is continuous and that observations vary continuously with θ. In particular, there exists a constant $K_\mathcal{D} > 0$ such that $|f_\theta(x)| \leq K_\mathcal{D}$ for every $x \in \mathcal{X}$ and $\theta \in \Theta$.

Dynamical models are inherently deterministic, as the sequence of observations generated by a model is fully determined once the initial condition is given. As noted in Section 1.2.1, each dynamical model has a set of invariant measures and these measures give rise to a family of stationary processes. In the large sample limit, fitting a family of dynamical models leads directly to a variational problem involving its associated processes. Conditions (D1)-(D3) ensure that the set of associated processes is non-empty and that the limiting variational problem is well-defined. Note, however, that we make no explicit assumptions about the invariant measures of any individual model.

In this paper our primary focus is on model families $\mathcal{D}$ having limited complexity. Roughly speaking, we require that the ℓ_p-covering numbers of the length-n real sequences produced by models in the family grow sub-exponentially in n at all scales. A precise notion of entropy for

model families is given in Section 1.4. In general, the elements of the family $\mathcal{D}$ may reflect chaotic behavior, or they may capture low order regularities such as periodicity, almost periodicity, or hierarchical structure. Examples and further discussion of model families can be found in Section 1.1. Fitting a family of such models to observations of a stochastic process amounts to looking for these types of regularities in the process.

Let $\mathcal{D}$ be a family of dynamical models that capture some behaviors of interest, and let $\mathbf{Y} = Y_0, Y_1, \ldots \in \mathbb{R}$ be an observed stationary ergodic process. Suppose that we wish to identify regularities in $\mathbf{Y}$ by fitting the observed values of the process with models in $\mathcal{D}$. Note that we do not assume that the observed process $\mathbf{Y}$ is generated by a process in $\mathcal{D}$, or that $\mathbf{Y}$ itself is of limited complexity. Let $\ell : \mathbb{R} \times \mathbb{R} \to \mathbb{R}$ be a nonnegative loss function that is jointly lower semicontinuous in its arguments. We require the following integrability condition:

$$\text{(C1)} \qquad \mathbb{E}\left[\sup_{|u| \le K_{\mathcal{D}}} \ell\big(u, Y_0\big)\right] \;<\; \infty.$$

(If the supremum in (C1) is not measurable, then one may replace the expectation by an outer expectation.) For each $n \ge 0$, $\theta \in \Theta$, and $x \in \mathcal{X}$ define

$$R_n(\theta : x) \;=\; \frac{1}{n} \sum_{k=0}^{n-1} \ell\big(f_\theta \circ T_\theta^k(x), Y_k\big)$$

to be the empirical risk of the model (T_θ, f_θ) with initial state x relative to the first n observations of $\mathbf{Y}$. We formalize empirical risk minimization as follows.

Definition 1.1. A sequence of measurable functions $\theta_n : \mathbb{R}^n \to \Theta$, $n \ge 1$, will be called (empirical) minimum risk estimates for $\mathcal{D}$ if

$$\lim_n \inf_x R_n(\hat{\theta}_n : x) \;=\; \lim_n \inf_\theta \inf_x R_n(\theta : x) \quad \text{w.p.1}, \tag{1.1}$$

where $\hat{\theta}_n := \theta_n(Y_0, \ldots, Y_{n-1})$.

Remark 1.2. Existence of the limit on the right hand side of (1.1) follows from Kingman's subadditive ergodic theorem (under (C1)). Existence of the limit on the left hand side of (1.1) is part of the definition.

Our primary goal is to investigate and characterize the limiting behavior of minimum risk estimates. Three main results are presented, corresponding to three levels of generality. At the highest level, we provide a variational characterization of the limiting behavior of minimum risk estimates (Theorem 1.12). We then focus on a natural signal plus

noise setting and show that if the family of dynamical models has low complexity then empirical risk minimization effectively separates the signal and the noise (Theorem 1.18). In the low complexity setting, we show that if the signal arises from a model in the family and the noise is appropriately centered with respect to the loss, then the minimum risk estimates are consistent (Theorem 1.21). A negative result (Proposition 1.22) shows that empirical risk minimization can be inconsistent when the family of dynamical models has high complexity.

Beyond the results above, the main contribution of our work is a systematic treatment of identifiability and complexity for families of dynamical models, with a focus on the misspecified case in which the observed process is not generated by a model in the family. We avoid assumptions on the invariant measures of the models in $\mathcal{D}$ by working with the processes they generate. In particular, the limit set of empirical risk estimators is characterized in terms of a distortion-based projection of the observed process onto the family of processes associated with $\mathcal{D}$. In addition, we introduce an entropy-based definition of complexity for families of dynamical models and families of infinite sequences that may be of independent interest. Our notion of entropy has close connections with topological entropy, studied in dynamical systems, and with stochastic mean widths, studied in empirical process theory.

Both the statements and proofs of our results rely on the concept of joinings, which are stationary couplings of stochastic processes. Joinings, introduced by Furstenberg [8], have been well-studied in ergodic theory, but have not been widely applied to problems of statistical inference. Our results show that joinings are intimately connected with minimum risk fitting of dynamical models. Several tools from the theory of joinings, including disjointness and relatively independent joinings, play an important role in our analysis.

1.1. **Examples.** Here we present several examples of families of dynamical models that capture potentially interesting types of regularities. Under some additional assumptions, these families have zero entropy, and therefore our general results apply to these models. In some cases, these or similar families have already been fit to data by applied scientists (e.g. [23, 41]), albeit without any theoretical guarantees of consistency.

Example 1.3. (Toral rotations and almost periodicity) Let the state space $\mathcal{X}$ be the d-dimensional torus $\mathbb{T}^d$, which is the direct product of d circles, $\mathbb{T}^d = S^1 \times \cdots \times S^1$. For a vector $\alpha \in \mathbb{T}^d$, define the transformation $R_\alpha : \mathbb{T}^d \to \mathbb{T}^d$ to be the rotation of $\mathbb{T}^d$ by the angle vector

α, i.e. $R_\alpha(x) = x + \alpha$ (addition in $\mathbb{T}^d$). Then let $\mathcal{F} \subset C(\mathbb{T}^d)$ be a compact set of continuous functions from $\mathbb{T}^d$ to $\mathbb{R}$ (with respect to the topology induced by the supremum norm). Let $\Theta = \mathbb{T}^d \times \mathcal{F}$, and define the family of dynamical models $\mathcal{D} = \{(R_\alpha, f) : (\alpha, f) \in \Theta\}$. With these definitions, $\mathcal{D}$ is a continuous family of dynamical models, and $h(\mathcal{D}) = 0$. Fitting this family to a process amounts to looking for periodic or "almost periodic" (also known as "quasi-periodic") structure in the observations. Intuitively, one is looking for up to d independent "periods" in a process.

As a specific example of a setting in which such models might arise, one may consider restricted classes of dynamic gene regulatory networks that exhibit periodic behavior. Inference of gene regulatory networks from observed data is considered an important problem in systems biology [21]. In recent years, it has become increasingly feasible for experimentalists to assay the abundance of all the genes in a given system with regular frequency over time. In such cases, one would like to infer the structure of the underlying network from the observed gene expression dynamics [23].

In many situations, one would expect gene regulatory networks to be zero entropy systems. In particular, the networks studied in chronobiology should exhibit periodic dynamics by definition. Examples of such systems include the cell cycle and circadian oscillators.

Example 1.4. (Subcritical logistic family and ecology) Since at least the early work of May [22], simple parametric families of dynamical systems have been used by ecologists as models of the population dynamics of many species [20]. In many instances, various types of deterministic models have been fit to ecological data (e.g. [41]).

The prototypical family in this context is the logistic family, which may be parametrized as follows. Consider the state space $\mathcal{X} = [0, 1]$ and the family of maps $T_a : [0, 1] \to [0, 1]$, where $T_a(x) = ax(1 - x)$ for $a \in [0, 4]$. If we restrict a to the region $[0, 3.5]$, then the family of dynamical models will have zero entropy. This situation is thought to occur in many naturally occurring populations (see results and discussion from [13]).

Example 1.5. (Symbolic dynamics and quasicrystals) In recent years, interest has grown in certain materials called quasicrystals. These materials were discovered in 1983 by Schechtman and exhibit an interesting mix of properties [36]. In particular, they have long-range aperiodic order (in contrast to crystals, which have long-range periodic order). Such systems often exhibit a rigid hierarchical structure at all scales, and they are frequently modeled using symbolic dynamical systems

[5, 35, 38], such as substitution systems. A full discussion of these systems is beyond the scope of this paper. However, if one were interested in looking for hierarchical structure in observational data, then these types of dynamical models could be used. Many families of such models naturally have zero entropy, as it is often forced by the long-range order and rigid hierarchical structure.

1.2. **Preliminaries.** In this section we introduce some concepts and notation that will be useful in what follows.

1.2.1. *Processes associated with dynamical models.* Let (T, f) be a dynamical model on a compact metrizable state space $\mathcal{X}$. Recall that a Borel probability measure μ on $\mathcal{X}$ is said to be *invariant* under T if $\mu(T^{-1}A) = \mu(A)$ for all Borel sets $A \subseteq \mathcal{X}$. Let $\mathcal{M}(\mathcal{X}, T)$ be the set of Borel measures on $\mathcal{X}$ that are invariant under T_θ, which is nonempty (see [45, p.152]). To each measure $\mu \in \mathcal{M}(\mathcal{X}, T)$ there is an associated real-valued process

$$\mathbf{U} = f(X), f(TX), f(T^2X), \ldots$$

where $X \in \mathcal{X}$ has distribution μ. The invariance of μ under T ensures that $\mathbf{U}$ is stationary. Here and in what follows we will regard real-valued processes as measures on the infinite product space $\mathbb{R}^{\mathbb{N}}$ equipped with its Borel sigma-field in the standard product topology.

Definition 1.6. Let $\mathcal{D} = \{(T_\theta, f_\theta) : \theta \in \Theta\}$ be a family of dynamical models. For each $\theta \in \Theta$ let

$$\mathcal{Q}_\theta \;=\; \Big\{\mathbf{U} = (f_\theta \circ T_\theta^k(X))_{k\geq 0} : X \sim \mu \text{ with } \mu \in \mathcal{M}(\mathcal{X}, T_\theta)\Big\}$$

be the set of stationary processes associated with (T_θ, f_θ), and let $\mathcal{Q}_\mathcal{D} = \bigcup_{\theta\in\Theta} \mathcal{Q}_\theta$ be the set of processes associated with the entire family of models $\mathcal{D}$.

1.2.2. *Joinings and distortion for stationary processes.* The statements and proofs of our principal results rely critically on stationary couplings of stationary processes, which are known as joinings.

Definition 1.7. A *joining* of two stationary processes $\mathbf{U} = \{U_k : k \geq 0\}$ and $\mathbf{V} = \{V_k : k \geq 0\}$ is a stationary process $\mathbf{W} = \{(\tilde{U}_k, \tilde{V}_k) : k \geq 0\}$ such that $\tilde{\mathbf{U}} = \{\tilde{U}_k : k \geq 0\}$ has the same distribution as $\mathbf{U}$ and $\tilde{\mathbf{V}} = \{\tilde{V}_k : k \geq 0\}$ has the same distribution as $\mathbf{V}$. Let $\mathcal{J}(\mathbf{U}, \mathbf{V})$ denote the family of all joinings of $\mathbf{U}$ and $\mathbf{V}$.

By definition, a joining of two stationary processes is a coupling of the processes that is itself stationary. Note that the family $\mathcal{J}(\mathbf{U}, \mathbf{V})$ always contains the the so-called independent joining under which $\tilde{\mathbf{U}}$

and $\tilde{\mathbf{V}}$ are independent copies of $\mathbf{U}$ and $\mathbf{V}$, respectively. Joinings were introduced by Furstenberg [8], and have been widely studied in ergodic theory [6, 9]. For notational convenience, we will frequently use $(\mathbf{U}, \mathbf{V})$ to denote a joining of $\mathbf{U}$ with $\mathbf{V}$. Also, note that the joining of three or more stationary processes may be defined analogously.

Definition 1.8. Let $L : \mathbb{R} \times \mathbb{R} \to \mathbb{R}$ be a nonnegative loss function. The L-distortion between two stationary processes $\mathbf{U}$ and $\mathbf{V}$ is given by

$$\gamma_L(\mathbf{U}, \mathbf{V}) \;=\; \inf_{\mathcal{J}(\mathbf{U},\mathbf{V})} \mathbb{E}\big[L(U_0, V_0)\big].$$

Remark 1.9. Joinings were used by Ornstein [30, 31, 32] to define the $\overline{d}$-distance between finite alphabet stationary processes based on the Hamming metric $\mathbb{I}(U_0 \neq V_0)$. The $\overline{d}$-distance was extended by Gray et al. [10] to stationary processes with general alphabets and to arbitrary metrics $\rho(U_0, V_0)$. The distortion $\gamma_L(\cdot, \cdot)$ is a straightforward generalization of these distances to nonnegative loss functions $L(\cdot, \cdot)$ that need not be metrics on $\mathbb{R}$.

Remark 1.10. The fact that the infimum defining $\gamma_L(\cdot, \cdot)$ runs over the set of joinings, rather than the set of couplings, is critical. A minimizing joining makes the average loss between elements of the process as small as possible over the entire future. By contrast, a minimizing coupling would make the processes as close as possible at time zero, without regard to their behavior in the future. Moreover, ergodic properties of the processes can severely constrain the set of possible joinings. For instance, for many pairs of processes $\mathbf{U}$ and $\mathbf{V}$, the only joining in $\mathcal{J}(\mathbf{U}, \mathbf{V})$ is the independent joining. Such processes are called disjoint.

1.3. **Convergence of Minimum Risk Estimates.** As noted above, a family $\mathcal{D}$ of dynamical models corresponds to a family $\mathcal{Q}_\mathcal{D} = \bigcup_{\theta\in\Theta} \mathcal{Q}_\theta$ of stationary processes. The problem of fitting models in $\mathcal{D}$ to an observed ergodic process $\mathbf{Y}$ using the loss $\ell(\cdot, \cdot)$ has a population analog in which we seek processes in $\mathcal{Q}_\mathcal{D}$ that minimize the distortion $\gamma_\ell(\cdot, \cdot)$ with $\mathbf{Y}$. The solution to the population problem is the γ_ℓ-projection of $\mathbf{Y}$ onto $\mathcal{Q}_\mathcal{D}$, and the corresponding set of parameters is a natural limit set for empirical risk estimators. This leads to the following definition, which is given for general loss functions.

Definition 1.11. Let $\mathcal{D}$ be a family of dynamical models parametrized by $\theta \in \Theta$. Given a nonnegative loss function $L : \mathbb{R} \times \mathbb{R} \to \mathbb{R}$ and a stationary ergodic process $\mathbf{Y}$, let

$$\Theta_L(\mathbf{Y}) \;=\; \operatorname{argmin}_{\theta\in\Theta} \min_{\mathbf{U}\in\mathcal{Q}_\theta} \gamma_L(\mathbf{U}, \mathbf{Y}),$$

which is the set of parameters θ such that some process in $\mathcal{Q}_\theta$ minimizes the distortion with $\mathbf{Y}$.

The proof of the following theorem, which relies on results of McGoff and Nobel [26], is presented in Section 2.

Theorem 1.12. *Let $\mathcal{D}$ be a family of dynamical models satisfying (D1)-(D3), and let $\ell(\cdot,\cdot)$ be a lower semicontinuous loss function. If $\mathbf{Y}$ is a stationary ergodic process satisfying (C1), then $\Theta_\ell(\mathbf{Y})$ is non-empty and compact and*

(1) *any sequence $\{\hat{\theta}_n = \theta_n(Y_0,\ldots,Y_{n-1})\}$ of minimum risk estimators converges almost surely to $\Theta_\ell(\mathbf{Y})$;*

(2) *for each $\theta \in \Theta_\ell(\mathbf{Y})$, there exists a sequence of minimum risk estimators that converges almost surely to θ.*

We emphasize that there is no assumed relationship between the observations $\mathbf{Y}$ and the family $\mathcal{D}$. Additionally, identifiability of parameters is addressed in a direct way, through the distortion γ_ℓ via the families $\mathcal{Q}_\theta$.

Theorem 1.12 shows that the limiting behavior of minimum risk estimators is characterized by the family $\Theta_\ell(\mathbf{Y})$, and in this way the theorem reduces the asymptotic analysis of empirical risk minimization to the analysis of this limit set. We show below how analysis of $\Theta_\ell(\mathbf{Y})$ yields both positive results (e.g. consistency) and negative results (inconsistency) in a signal plus noise setting.

1.4. **Entropy for sequence families.** A key issue in nonparametric inference is how to assess the complexity of a family of models. Complexity measures play an important role in establishing consistency, convergence rates, and optimality for a variety of inference procedures. Although fitting nonlinear dynamical models to ergodic processes is substantially different from model fitting for classification or regression, complexity still plays an important role in the consistency of empirical risk minimization.

We assess the complexity of a family $\mathcal{D}$ through the covering numbers of the infinite real-valued sequences generated by its constituent models. From a statistical point of view, it is natural to consider empirical ℓ_2 covering numbers, while from a dynamical systems point of view, it is natural to consider empirical ℓ_∞ covering numbers, as is done with topological entropy [45]. As we show below, these two approaches coincide.

Let $\mathbf{u} = (u_k)_{k\geq 0}$ and $\mathbf{v} = (v_k)_{k\geq 0}$ denote infinite sequences in $\mathbb{R}^{\mathbb{N}}$. For each $n \geq 1$ and $1 \leq p \leq \infty$, define pseudo-metrics $d_{n,p}(\cdot,\cdot)$ as follows:

$$d_{n,p}(\mathbf{u},\mathbf{v}) = \begin{cases} \left(n^{-1}\sum_{k=0}^{n-1} |u_k - v_k|^p\right)^{1/p} & \text{if } 1 \leq p < \infty \\ \max_{0\leq k\leq n-1} |u_k - v_k| & \text{if } p = \infty. \end{cases}$$

Let $\mathcal{U} \subseteq \mathbb{R}^{\mathbb{N}}$ be a family of infinite sequences. For each $r > 0$ let $N(\mathcal{U}, r, d_{n,p})$ denote the covering number of the set $\mathcal{U}$ under the pseudo-metric $d_{n,p}(\cdot,\cdot)$ at radius r. Let

$$h_p(\mathcal{U}, r) \;=\; \limsup_n \frac{1}{n} \log N(\mathcal{U}, r, d_{n,p}),$$

which is the exponential growth rate of the covering numbers at radius r, and define the ℓ_p entropy of the family $\mathcal{U}$ as the supremum of these growth rates, namely

$$h_p(\mathcal{U}) \;=\; \lim_{r\searrow 0} h_p(\mathcal{U}, r).$$

The following result is established in Section 3.

Theorem 1.13. *The ℓ_p entropies $h_p(\mathcal{U})$ for $1 \leq p \leq \infty$ are all equal.*

Remark 1.14. Although it is not needed here, we note that Theorem 1.13 holds more generally for sets of sequences $\mathcal{U} \subseteq A^{\mathbb{N}}$ where (A, ρ) is any metric space such that

$$\lim_{r\searrow 0} r \log N(A, r, \rho) = 0$$

and the pseudo metrics $d_{n,p}(\cdot,\cdot)$ are defined in terms of ρ.

Definition 1.15 (Entropy of a dynamical family)**.** The entropy $h(\mathcal{D})$ of a family $\mathcal{D}$ of dynamical models is the common value of $h_p(\mathcal{U}_\mathcal{D})$, where

$$\mathcal{U}_\mathcal{D} = \{(f_\theta \circ T_\theta^k(x))_{k\geq 0} : x \in \mathcal{X},\, \theta \in \Theta\} \;\subseteq\; \mathbb{R}^{\mathbb{N}} \tag{1.2}$$

is the set of infinite sequences generated by models in $\mathcal{D}$.

Remark 1.16. It is straightforward to show that $\mathcal{U}_\mathcal{D}$ is a compact subset of $\mathbb{R}^{\mathbb{N}}$ in its product topology. Let $\tau : \mathbb{R}^{\mathbb{N}} \to \mathbb{R}^{\mathbb{N}}$ be the left-shift map defined by $\tau(\mathbf{u})_k = u_{k+1}$ for $k \geq 0$. Then it is easy to see that τ is continuous and $\tau(\mathcal{U}_\mathcal{D}) \subset \mathcal{U}_\mathcal{D}$. Thus $(\mathcal{U}_\mathcal{D}, \tau)$ is a topological dynamical system that captures the dynamics of the family $\mathcal{D}$, and the entropy $h(\mathcal{D})$ defined above is the topological entropy of this system.

We note that the entropy $h(\mathcal{D})$ may also be characterized in terms of the entropies $h(\mathbf{U})$ of the processes in $\mathbf{U} \in \mathcal{Q}_{\mathcal{D}}$, which we define in Section 3.3. The following lemma is established in Appendix A.

Lemma 1.17. *For any family $\mathcal{D}$ of dynamical models satisfying (D1)-(D3),*

$$h(\mathcal{D}) = \sup_{\mathbf{U} \in \mathcal{Q}_{\mathcal{D}}} h(\mathbf{U}).$$

1.5. **Signal Plus Noise.** We now turn our attention to empirical risk minimization when the observed process $\mathbf{Y}$ has a signal plus noise structure. We assume in what follows that $Y_k = V_k + \varepsilon_k$ for each $k \geq 0$, where $\mathbf{V} = \{V_k : k \geq 0\}$ is a stationary ergodic process and $\boldsymbol{\varepsilon} = \{\varepsilon_k : k \geq 0\}$ is an i.i.d. noise process that is independent of $\mathbf{V}$. We indicate this relationship using the process-sum notation $\mathbf{Y} = \mathbf{V} + \boldsymbol{\varepsilon}$. We require the following integrability conditions:

$$\text{(C1)} \qquad \mathbb{E}\left[\sup_{|x| \leq K_{\mathcal{D}}} \ell(x, Y_0)\right] < \infty;$$

$$\text{(C2)} \qquad \mathbb{E}\left[\sup_{|x| \leq K_{\mathcal{D}}} \ell(x, V_0)\right] < \infty;$$

$$\text{(C3)} \qquad \text{for all } u, v \in \mathbb{R}, \quad \mathbb{E}\,\ell(u, v + \varepsilon_0) < \infty.$$

Note that (C1) is the same condition required in the general setting, and (C2) and (C3) refer only to $\mathbf{V}$ and $\boldsymbol{\varepsilon}$, respectively.

Theorem 1.12 ensures that any sequence of minimum risk estimators will converge to the set $\Theta_\ell(\mathbf{Y})$ of optimal parameters for $\mathbf{Y}$. Of interest here is when and whether empirical risk minimization can decouple the signal from the noise and recover the optimal parameters $\Theta_\ell(\mathbf{V})$ for the signal process $\mathbf{V}$. We begin with the following general result.

Theorem 1.18. *Let $\mathbf{Y} = \mathbf{V} + \boldsymbol{\varepsilon}$ satisfy (C1)-(C3), and let $\mathcal{D}$ satisfy (D1)-(D3). If $h(\mathcal{D}) = 0$, then any sequence of minimum ℓ-risk estimates converges almost surely to $\Theta_L(\mathbf{V})$, where $L(u, v) := \mathbb{E}\,\ell(u, v + \varepsilon_0)$.*

Remark 1.19. Since $\ell(\cdot, \cdot)$ is nonnegative and lower semicontinuous, the auxiliary loss function $L(\cdot, \cdot)$ has the same properties (using Fatou's Lemma for the lower semicontinuity).

Remark 1.20. If a given process $\mathbf{Y}$ can be expressed in two different ways as $\mathbf{Y} = \mathbf{V} + \boldsymbol{\varepsilon}$ and $\mathbf{Y} = \mathbf{V}' + \boldsymbol{\varepsilon}'$, then the proof of Theorem 1.18 shows that $\Theta_L(\mathbf{V}) = \Theta_{L'}(\mathbf{V}')$, where L' is defined using $\boldsymbol{\varepsilon}'$ in place of $\boldsymbol{\varepsilon}$.

Theorem 1.18 shows that if $h(\mathcal{D}) = 0$, then empirical risk minimization does indeed decouple the signal from the noise. However, the presence of the auxiliary loss function $L(\cdot, \cdot)$ in the conclusion of the theorem begs the question of whether the limit set is equal to the limiting parameter set $\Theta_\ell(\mathbf{V})$ associated with the signal, as one would like. The following two results address this question under additional hypotheses.

Theorem 1.21. *Let $\mathbf{Y} = \mathbf{V} + \boldsymbol{\varepsilon}$ satisfy (C1)-(C3), and let $\mathcal{D}$ be a family of dynamical models satisfying (D1)-(D3) with $h(\mathcal{D}) = 0$. Let $\{\hat{\theta}_n : n \geq 1\}$ be any sequence of minimum risk estimators based on $\mathbf{Y}$.*

(1) *If $\ell(u, v) = D_F(v, u)$ is a Bregman divergence and $\mathbb{E}\,\varepsilon_0 = 0$, then $\hat{\theta}_n$ converges almost surely to $\Theta_\ell(\mathbf{V})$.*

(2) *If $\mathbf{V}$ is an ergodic process in $\mathcal{Q}_{\theta_0}$ and $\mathbb{E}\,\ell(u, v + \varepsilon_0) \geq \mathbb{E}\,\ell(0, \varepsilon_0)$ for all u, v, with equality if and only if $u = v$, then $\hat{\theta}_n$ converges almost surely to $\{\theta \in \Theta : \mathbf{V} \in \mathcal{Q}_\theta\}$.*

Theorem 1.21 establishes that minimum risk fitting of a zero-entropy dynamical family effectively isolates the signal under two types of hypotheses. The first places restrictions on the loss function but allows the observation process to be quite general. The second places restrictions on the signal and on the joint behavior of the loss function and the noise processes.

Without the entropy condition $h(\mathcal{D}) = 0$, the conclusion of Theorem 1.21 does not hold in general, as the following proposition highlights. For $\theta \in \Theta$, let $\mathcal{D}_\theta = \{(T_\theta, f_\theta)\}$ be the family consisting of the single dynamical model associated to θ.

Proposition 1.22. *Let $\mathcal{D}$ be a family of dynamical models satisfying (D1)-(D3). Suppose that*

- *$h(\mathcal{D}) > 0$;*
- *there exists θ_0 such that $h(\mathcal{D}_{\theta_0}) = 0$ and $\mathcal{Q}_{\theta_0} \setminus \cup_{\theta \neq \theta_0} \mathcal{Q}_\theta$ contains an ergodic process $\mathbf{V}$.*

Then there exists $\sigma_0 > 0$ such that for every i.i.d. process $\boldsymbol{\varepsilon}$ with $\varepsilon_i \sim N(0, \sigma^2)$ for $\sigma > \sigma_0$, the least squares estimates derived from $\mathbf{Y} = \mathbf{V} + \boldsymbol{\varepsilon}$ converge to a compact set Θ_* such that $\mathbf{V} \notin \cup_{\theta \in \Theta_*} \mathcal{Q}_\theta$. Moreover, such a family $\mathcal{D}$ exists.

Under the conditions of this proposition, inconsistency holds almost surely despite the fact that the signal process $\mathbf{V}$ is generated by a dynamical model in the family. The idea underlying this phenomenon

is that since $\mathcal{D}$ has positive entropy, it is capable of tracking the noise, and then the least squares estimates will overfit the observed sequence.

1.6. **Squared loss and mean width.** In this section we give special attention to the squared loss, $\ell(x,y) = (x-y)^2$. Analysis of the squared loss naturally leads to another measure of complexity of the family $\mathcal{D}$ based on the mean width of sequences generated by $\mathcal{D}$ with respect to the noise.

Definition 1.23. Let $\mathcal{D}$ be a family of dynamical models, and let $\boldsymbol{\varepsilon} = (\varepsilon_k)_{k\geq 0}$ be an i.i.d. process with mean zero and finite variance. The n-sample mean width of $\mathcal{D}$ relative to $\boldsymbol{\varepsilon}$ is

$$\kappa_n(\mathcal{D} : \boldsymbol{\varepsilon}) \;=\; \mathbb{E}\left[\sup_{x,\theta} \sum_{k=0}^{n-1} f_\theta \circ T_\theta^k(x) \cdot \varepsilon_k\right]. \tag{1.3}$$

Define the mean width of $\mathcal{D}$ relative to $\boldsymbol{\varepsilon}$ to be the limiting linear growth rate of the finite sample mean widths,

$$\kappa(\mathcal{D} : \boldsymbol{\varepsilon}) = \lim_n \frac{1}{n}\kappa_n(\mathcal{D} : \boldsymbol{\varepsilon}), \tag{1.4}$$

which exists by subadditivity (see Remark 3.2). When $\varepsilon_i \sim N(0,1)$ we write $\kappa(\mathcal{D} : \boldsymbol{\varepsilon})$ as $\kappa_G(\mathcal{D})$ and refer to this quantity as the Gaussian mean width of the family $\mathcal{D}$.

Finite sample mean widths have been widely studied in machine learning and empirical process theory, with an emphasis on Rademacher and Gaussian noise [3, 18]. The mean width of $\mathcal{D}$ has close connections with the entropy of $\mathcal{D}$.

Theorem 1.24. *Let $\boldsymbol{\varepsilon} = (\epsilon_k)_{k\geq 0}$ be i.i.d. with mean zero and finite variance. If $h(\mathcal{D}) = 0$ then $\kappa(\mathcal{D} : \boldsymbol{\varepsilon}) = 0$. Moreover, the Gaussian mean width $\kappa_G(\mathcal{D}) = 0$ if and only if $h(\mathcal{D}) = 0$.*

Remark 1.25. In general, one cannot expect a more quantitative relationship between asymptotic mean width and entropy. Indeed, the entropy $h(\mathcal{D})$ is invariant under any continuous, invertible change of coordinates (see [45, p. 167]), whereas the mean width $\kappa_G(\mathcal{D})$ is not invariant under such operations. In particular, the asymptotic mean width scales linearly if the observation functions f_θ are multiplied by a fixed constant, while the entropy remains unchanged.

Let γ_2 denote the ℓ-distortion in the special case that ℓ is the squared loss. Note that $\gamma_2(\mathbf{U},\mathbf{V})^{1/2}$ is in fact a metric on the space of $\mathbb{R}$-valued stationary stochastic processes. As the squared loss is a Bregman divergence, minimum risk fitting of a zero entropy family will converge

to the optimal parameter set for the signal by Theorem 1.21. The next theorem extends this result to the case where the mean width of the family is zero.

Theorem 1.26. *Let* $\mathbf{Y} = \mathbf{V} + \boldsymbol{\varepsilon}$*, where* $\mathbf{V}$ *is ergodic and* $\boldsymbol{\varepsilon}$ *is an i.i.d. process with mean zero and finite variance. If* $\kappa(\mathcal{D} : \boldsymbol{\varepsilon}) = 0$*, then any sequence of least squares estimators converges almost surely to* $\Theta_2(\mathbf{V})$*.*

In our final result of this section, we establish the consistency of least squares estimation for a family of transformations on a compact state space in $\mathbb{R}^d$ where each observation function is the identity. Suppose $\mathcal{X} \subset \mathbb{R}^d$ is compact and $\{T_\theta : \theta \in \Theta\}$ is a family of transformations on $\mathcal{X}$ such that Θ is a compact metric space and $(\theta, x) \mapsto T_\theta(x)$ is continuous. Further, suppose that $Y_k = T^k_{\theta^*}(X) + \varepsilon_k$ where X is distributed according to an ergodic measure $\mu \in \mathcal{M}(\mathcal{X}, T_{\theta^*})$, and $\boldsymbol{\varepsilon} = (\varepsilon_k)_{k \geq 0}$ is i.i.d. with mean zero and finite variance and is independent of X.

Corollary 1.27. *If the topological entropy of* T_θ *is zero for all* $\theta \in \Theta$*, then any sequence of least squares estimators converges almost surely to the set* $\{\theta \in \Theta : \mu(T_\theta = T_{\theta^*}) = 1\}$*.*

The limit set in Corollary 1.27 contains θ^* and serves as the natural identifiability class of θ^* in this setting.

1.7. **Related work.** Following the development of the statistical learning theory for i.i.d. data by Vapnik and Chervonenkis [43], there has been substantial interest in learning from dependent data (e.g. [1, 2, 11, 14, 28, 47]). In many such cases, the authors consider learning from dependent processes satisfying some type of mixing condition (see [4]), and the results yield learning rates for prediction. Other results provide sufficient (and sometimes necessary) conditions for the consistency of certain prediction and classification procedures from dependent processes [27, 29]. The present work is concerned with parameter or feature estimation, which is related to yet distinct from the problems considered in the previous works.

The work of Ornstein and Weiss [33] studies the estimation of a stochastic process from its samples. In that work, the authors propose a particular inference procedure and characterize when it produces consistent estimates of the stochastic process in the d-bar metric. In our work, we focus on a different procedure, based on empirical risk minimization as opposed to matching k-block frequencies, and we consider consistency in the context of restricted families of models, rather than in the class of all stationary ergodic processes.

Our work is also somewhat related to a line of research concerning least squares estimation (LSE) of individual sequences from noisy observations, which has been considered in a variety of contexts by several authors (see [34, 42, 46], among others). The most recent of these works ([34]) employs the tools of empirical process theory to prove results on both consistency and asymptotic normality of LSE of individual sequences from signal plus noise. In the present work, we consider only specific sets of individual sequences: those that arise from a continuous family of dynamical models, as in (1.2), and we are interested in inference of a dynamical invariant parameter (i.e. θ), rather than the signal sequence itself. The dynamical invariance of our model sequences, along with the invariance of our inference target θ, allows us to obtain consistency of LSE under our general mean width (or entropy) conditions, rather than under the more restrictive complexity conditions appearing in [34]. Furthermore, our setting distinguishes itself from this related work by naturally handling a substantial degree of model misspecification.

In Furstenberg's original work on joinings [8], he includes an application of joinings to a nonlinear filtering problem. Beyond this application, we are not aware of other uses of joinings in the literature on statistical inference.

In recent work of Lev, Peled, and Peres [19], some of Furstenberg's original results have been substantially extended. This work concerns two problems: given an infinite sequence of signal plus noise, first detect whether the signal is non-zero, and second recover the signal from the given sequence. Although somewhat related to our work, these results possess significant differences from ours. For one thing, the detection and filtering problems mentioned above place no restrictions beyond measurability on the detection and filtering mappings, whereas we specifically consider finitary inference via empirical risk minimization. Additionally, they consider arbitrary sets of individual sequences with proper specification (as in [34]), where as we consider potentially misspecified parameter estimation of dynamical models.

Finally, we mention that many questions about statistical inference in the context of dynamical systems have been considered in a wide variety of subject areas; see the survey [25] for a broad overview and references. Let us mention that dynamical systems in the observational noise setting have been studied in [15, 16, 24], and statistical prediction in the context of dynamical systems has been considered in [11, 12, 39, 44].

1.8. **Generalizations.** Generalization of all of the definitions and results of the paper to $\mathbb{R}^d$-valued models and processes is straightforward, requiring only minor changes of notation. We omit the details. In a different direction, one could analyze families of dynamical models defined on a non-compact state space $\mathcal{X}$ with uniformly bounded observation functions, requiring only measurability of the maps $(\theta, x) \mapsto T_\theta(x)$ and $(\theta, x) \mapsto f_\theta(x)$. For families $\mathcal{D}$ of this more general type, the set $\mathcal{U}_\mathcal{D}$ of associated sequences would not necessarily be a closed (hence compact) subset of $\mathbb{R}^\mathbb{N}$, and in this case one needs to consider the closure of $\mathcal{U}_\mathcal{D}$, along with all the stationary processes supported on this set. The analysis here can be carried out in this more general setting, but the corresponding results are difficult to interpret in the context of the original inference problem.

1.9. **Outline of the rest of the paper.** Section 2 contains the proof of the general convergence theorem. Proofs of the results on entropy and mean width appear in Section 3. The main inference results for the signal plus noise setting are established in Section 4. Section 5 contains the proof of the negative result Proposition 1.22.

2. Optimal tracking and proof of Theorem 1.12

In this section we discuss the connections between fitting dynamical models and the optimal tracking problem studied in [26]. In particular, we construct a single dynamical system that captures the important features of the family $\mathcal{D}$ of dynamical models, and we show how this system may be analyzed in the context of optimal tracking.

2.1. **Optimal tracking.** The tracking problem for dynamical systems concerns two systems: a model system $T : \mathcal{Z} \to \mathcal{Z}$, where $\mathcal{Z}$ is compact and metrizable and T is continuous, and an an observed system $S : \mathcal{Y} \to \mathcal{Y}$, where $\mathcal{Y}$ is a separable completely metrizable space and S is Borel measurable. Given an initial segment of a trajectory $y, S(y), \ldots, S^{n-1}(y)$ from the observed system, one seeks a corresponding initial condition z_n such that the trajectory $z_n, T(z_n), \ldots, T^{n-1}(z_n)$ from the model system "tracks" the given trajectory from the observed system. An optimal tracking trajectory is chosen by minimizing an additive cost functional

$$\sum_{k=0}^{n-1} c\big(S^k y, T^k z\big),$$

where $c : \mathcal{Y} \times \mathcal{Z} \to \mathbb{R}$ is a fixed lower semicontinuous cost function.

The results from [26] consider the situation when the observed initial condition y is drawn from an ergodic measure $\nu \in \mathcal{M}(\mathcal{Y}, S)$ and $\sup_z |c(y,z)|$ is bounded above by a function in $L^1(\nu)$. Here we state a version of the previous results that is sufficient for our purposes, for which we require a bit more notation. Let Θ be a compact metrizable space, and let $\varphi : \mathcal{Z} \to \Theta$ be a continuous map satisfying $\varphi \circ T = \varphi$. We denote by $\mathcal{J}(\nu : \theta)$ the set of joinings of the process $\{S^k(Y_0)\}_{k\geq 0}$, where $Y_0 \sim \nu$, with any process of the form $\{T^k(Z_0)\}_{k\geq 0}$, where $Z_0 \sim \mu$ for some $\mu \in \mathcal{M}(\mathcal{Z}, T)$ such that $\mu(\varphi^{-1}\{\theta\}) = 1$.

Theorem A ([26])**.** *Let $T : \mathcal{Z} \to \mathcal{Z}$, $S : \mathcal{Y} \to \mathcal{Y}$, $c : \mathcal{Y} \times \mathcal{Z} \to \mathbb{R}$, ν, Θ, and $\varphi : \mathcal{Z} \to \Theta$ be as above. If $\hat{z}_n = \hat{z}_n(y, \ldots, S^{n-1}y)$ and the following equality holds ν almost surely,*

$$\lim_n \frac{1}{n} \sum_{k=0}^{n-1} c(S^k y, T^k \hat{z}_n) = \lim_n \inf_{z \in \mathcal{Z}} \frac{1}{n} \sum_{k=0}^{n-1} c(S^k y, T^k z), \tag{2.1}$$

then $\hat{\theta}_n = \varphi(\hat{z}_n)$ converges (ν almost surely) to the non-empty, compact set

$$\Theta_{min} = \operatorname{argmin}_{\theta \in \Theta} \min_{\mathcal{J}(\nu:\theta)} \mathbb{E}\big[c(Y_0, Z_0)\big]. \tag{2.2}$$

Furthermore, for any $\theta \in \Theta_{min}$, there exists $\hat{z}_n$ such that (2.1) holds and $\hat{\theta}_n = \varphi(\hat{z}_n)$ converges to θ.

2.2. **Proof of Theorem 1.12.** To be begin the proof, we describe how fitting a continuous family of dynamical models to an observed stochastic process can be cast as a tracking problem. As an important first step, we define a single dynamical system that encapsulates the entire family of dynamical models. Consider the state space

$$\mathcal{Z} = \left\{ \left(\theta,\, \left(f_\theta \circ T_\theta^k(x)\right)_{k\geq 0}\right) : \theta \in \Theta,\, x \in \mathcal{X} \right\} \subseteq \Theta \times \mathbb{R}^{\mathbb{N}},$$

and define the transformation $T : \mathcal{Z} \to \mathcal{Z}$ by $T(\theta, (u_k)_{k\geq 0}) = (\theta, (u_{k+1})_{k\geq 0})$, which is clearly continuous (in the product topology). We now establish some basic properties of the dynamical system $(\mathcal{Z}, T)$ associated with the family $\mathcal{D}$ of dynamical models.

Lemma 2.1. *The set $\mathcal{Z}$ is a compact subset of $\Theta \times \mathbb{R}^{\mathbb{N}}$ when $\mathbb{R}^{\mathbb{N}}$ is equipped with the product topology, and the map T is continuous. If μ is an ergodic element of $\mathcal{M}(\mathcal{Z}, T)$, then there exists $\theta \in \Theta$ and an ergodic process $\mathbf{U} \in \mathcal{Q}_\theta$ with distribution ν such that $\mu = \delta_\theta \otimes \nu$.*

Proof. By our hypotheses on the family $\mathcal{D}$, both the parameter space Θ and the state space $\mathcal{X}$ are compact, and therefore $\Theta \times \mathcal{X}$ is compact.

Define the map $\pi : \Theta \times \mathcal{X} \to \Theta \times \mathbb{R}^{\mathbb{N}}$ by

$$\pi(\theta, x) = \Big(\theta, \big(f_\theta \circ T_\theta^k(x)\big)_{k\geq 0}\Big).$$

It is clear from the definition of $\mathcal{Z}$ that $\mathcal{Z}$ is the image of $\Theta \times \mathcal{X}$ under π. To show that $\mathcal{Z}$ is compact, it now suffices to check that π is continuous.

Let $\{(\theta_n, x_n)\}_{n\geq 1}$ be a sequence converging to (θ, x) in $\Theta \times \mathcal{X}$. Let $K \in \mathbb{N}$. The continuity conditions (D2) and (D3) imply that for $0 \leq k \leq K$,

$$\lim_n f_{\theta_n} \circ T_{\theta_n}^k(x_n) = f_\theta \circ T_\theta^k(x).$$

As K was arbitrary, we have shown that $\{\pi(\theta_n, x_n)\}_{n\geq 1}$ converges to $\pi(\theta, x)$ in $\Theta \times \mathbb{R}^{\mathbb{N}}$ in the product topology, and therefore π is continuous.

The left-shift $\tau : \mathbb{R}^{\mathbb{N}} \to \mathbb{R}^{\mathbb{N}}$ is continuous, and therefore $T = (\mathrm{Id} \times \tau)|_{\mathcal{Z}}$ is continuous.

For the last statement of the lemma, define the map $R : \Theta \times \mathcal{X} \to \Theta \times \mathcal{X}$ as the skew-product over the identity, $R(\theta, x) = (\theta, T_\theta(x))$. By construction, we have $\pi \circ R = T \circ \pi$. Thus, π is a factor map from $(\Theta \times \mathcal{X}, T)$ onto $(\mathcal{Z}, S)$. It follows that the push-forward map from $\mathcal{M}(\Theta \times \mathcal{X}, R)$ to $\mathcal{M}(\mathcal{Z}, T)$ (given by $\eta \mapsto \eta \circ \pi^{-1}$) is a surjection [7, p. 19].

Now let $\mu \in \mathcal{M}(\mathcal{Z}, T)$ be ergodic. Since the push-forward map from $\mathcal{M}(\Theta \times \mathcal{X}, R)$ to $\mathcal{M}(\mathcal{Z}, T)$ (given by $\eta \mapsto \eta \circ \pi^{-1}$) is a surjection, there exists an ergodic $\eta \in \mathcal{M}(\Theta \times \mathcal{X}, R)$ such that $\eta \circ \pi^{-1} = \mu$. Since $\mathrm{proj}_\Theta \circ T = \mathrm{proj}_\Theta$, the induced measure $\eta \circ (\mathrm{proj}_\Theta \circ \pi)^{-1}$ on Θ must be invariant under the identity map. Also, it must be ergodic, since η is ergodic. As the only ergodic measures for the identity map are the point masses, we see that there exists $\theta \in \Theta$ such that $\eta \circ (\mathrm{proj}_\Theta \circ \pi)^{-1} = \delta_\theta$. Then $\eta = \delta_\theta \otimes \xi$ for some ergodic measure $\xi \in \mathcal{M}(\mathcal{X}, T_\theta)$. Finally, we conclude that $\mu = \eta \circ \pi^{-1} = \delta_\theta \otimes \nu$, where ν is the distribution of an ergodic process in $\mathcal{Q}_\theta$. □

We now proceed with the proof of Theorem 1.12. The observed process $\mathbf{Y}$ gives rise to an observed dynamical system in the tracking problem, where $\mathcal{Y} = \mathbb{R}^{\mathbb{N}}$, $S : \mathcal{Y} \to \mathcal{Y}$ is the left shift $S((u_k)_{k\geq 0}) = (u_{k+1})_{k\geq 0}$, and ν is the distribution of $\mathbf{Y}$ on $\mathbb{R}^{\mathbb{N}}$. Finally, as a cost function, we choose $c : \mathcal{Y} \times \mathcal{Z} \to \mathbb{R}$, where

$$c\Big(\mathbf{v}, \big(\theta, \mathbf{u}\big)\Big) = \ell(u_0, v_0)$$

Then an application of Theorem A yields that any sequence of minimum ℓ-risk parameters $(\hat{\theta}_n)_{n\geq 1}$ converges almost surely to the set

$\Theta_\ell(\mathbf{Y})$. Furthermore, this projection is nonempty and compact, and Theorem 1.12 (2) holds. We have thus proved Theorem 1.12.

3. Entropy and mean width

In this section we study the notions of entropy and mean width for families of dynamical models. We begin with entropy.

3.1. **Entropy for families of dynamical models.** Recall that the pseudo-metrics $d_{n,p}(\cdot,\cdot)$ and the ℓ_p entropies $h_p(\mathcal{D})$ were defined in Section 1.4. Additionally, we define $B_{n,p}(\mathbf{u},r)$ to be the $d_{n,p}(\cdot,\cdot)$-ball of radius r centered at $\mathbf{u}\in\mathbb{R}^{\mathbb{N}}$.

Before proving the main results presented in Section 1.4, we make a simple observation. For $p\geq 1$, we have that $d_{n,p}(\mathbf{u},\mathbf{v})\leq d_{n,\infty}(\mathbf{u},\mathbf{v})$. Hence, for any $\mathcal{U}\subset\mathbb{R}^{\mathbb{N}}$ and $\delta>0$,

$$N(\mathcal{U},\delta,d_{n,p})\leq N(\mathcal{U},\delta,d_{n,\infty}). \tag{3.1}$$

The following lemma, which bounds the cardinality of $d_{n,\infty}(\cdot,\cdot)$-separated sets that are contained in a single $d_{n,p}(\cdot,\cdot)$-ball, is used to prove Theorem 1.13. For notation, let $M_n^\infty(\mathcal{U},\delta)$ denote the maximum cardinality of any set in $\mathcal{U}$ that is δ-separated with respect to $d_{n,\infty}(\cdot,\cdot)$.

Lemma 3.1. *Let $p\in[1,\infty)$, $K\geq 1$ and $\delta\in(0,1)$. Set $\epsilon=(\delta/2)^{(1+p)/p}$. Then for any $\mathbf{u}\in[-K,K]^{\mathbb{N}}$ and $n\geq 1$,*

$$M_n^\infty(B_{n,p}(\mathbf{u},\epsilon),\delta)\leq(3K/\delta)^{\delta n/2}\cdot 2^{H(\delta/2)n},$$

where $H(x)=-x\log x-(1-x)\log(1-x)$.

Proof. Suppose $\mathbf{u}\in[-K,K]^{\mathbb{N}}$ and there exists $\{\mathbf{v}_1,\dots,\mathbf{v}_M\}\subset B_{n,p}(\mathbf{u},\epsilon)$ such that $d_{n,\infty}(\mathbf{v}_i,\mathbf{v}_j)\geq\delta$ for $i\neq j$. For the sake of notation, let $\mathbf{u}=(u(k))_{k\geq 0}$ and $\mathbf{v}_j=(v_j(k))_{k\geq 0}$. By subtracting $\mathbf{u}$ from all these sequences if necessary, we assume without loss of generality that $u(k)=0$ for all k.

The idea of the proof is to bound M by estimating the number of coordinates of each $\mathbf{v}_j$ that deviate from 0 by more than $\delta/2$. To begin, we define the following subsets of $[-K,K]$: $A_0=[-\delta/2,\delta/2)$, and $A_r=[-K+(r-1)\delta,-K+r\delta)\setminus A_0$, for $r=1,\dots,s$, where $s=\lceil 2K/\delta\rceil$. Now we code the points $\mathbf{v}_j$ according to this partition: define $r_j(k)$ by the relation $v_j(k)\in A_{r_j(k)}$, and let $\pi:\{\mathbf{v}_1,\dots,\mathbf{v}_M\}\to\{0,\dots,s\}^n$ be given by

$$\pi(\mathbf{v}_j)=(r_j(0),\dots,r_j(n-1)).$$

First, observe that π is injective. Indeed, if $i\neq j$, then $d_n^\infty(\mathbf{v}_i,\mathbf{v}_j)\geq\delta$, and hence there exists k such that $|v_i(k)-v_j(k)|\geq\delta$, which implies

$r_i(k) \neq r_j(k)$. Second, observe that

$$s = \lceil 2K/\delta \rceil \leq 2K/\delta + 1 \leq 3K/\delta. \tag{3.2}$$

Now we proceed with the main bounds. Since $\{\mathbf{v}_1, \ldots, \mathbf{v}_M\} \subset B_{n,p}(\mathbf{u}, \epsilon)$, we have that for each j,

$$\epsilon^p n \geq \sum_{k=1}^{n} |v_j(k)|^p.$$

Furthermore, by construction, if $r_j(k) \neq 0$, then $|v_j(k)| \geq \delta/2$, and therefore

$$\epsilon^p n \geq \sum_{k=1}^{n} |v_j(k)|^p \geq (\delta/2)^p \big|\{k : r_j(k) \neq 0\}\big|.$$

From this inequality and the choice of ϵ, we deduce that

$$\{\pi(\mathbf{v}_1), \ldots, \pi(\mathbf{v}_M)\} \subset \left\{ z \in \{0, \ldots, s\}^n : \big|\{k : z(k) \neq 0\}\big| \leq \delta n/2 \right\}. \tag{3.3}$$

By the facts established above (injectivity of π, (3.3), and (3.2)) and a well-known inequality for binomial sums (see [37]),

$$\begin{aligned} M = \big|\{\pi(\mathbf{v}_1), \ldots, \pi(\mathbf{v}_M)\}\big| &\leq \sum_{k=0}^{\delta n/2} \binom{n}{k} s^k \leq s^{\delta n/2} \sum_{k=0}^{\delta n/2} \binom{n}{k} \\ &\leq (3K/\delta)^{\delta n/2} \sum_{k=0}^{\delta n/2} \binom{n}{k} \leq (3K/\delta)^{\delta n/2} \cdot 2^{H(\delta/2)n}, \end{aligned}$$

which completes the proof of the lemma. □

With this lemma in place, we now turn to the proof of Theorem 1.13.

Proof of Theorem 1.13. The inequality $h_p(\mathcal{D}) \leq h_\infty(\mathcal{D})$ follows easily from (3.1) and the definition of entropy. The remainder of the proof is devoted to showing the reverse inequality. As in the definition of entropy, we let

$$\mathcal{U} = \left\{ \big(f_\theta \circ T_\theta^k(x)\big)_{k \geq 0} : x \in \mathcal{X},\, \theta \in \Theta \right\}.$$

Since $\mathcal{X}$ and Θ are compact and the map $(\theta, x) \to f_\theta(x)$ is continuous, there exists $K \geq 1$ such that $|f_\theta(x)| \leq K$ for all $x \in \mathcal{X}$ and $\theta \in \Theta$. Thus $\mathcal{U} \subset [-K, K]^{\mathbb{N}}$.

Let $\delta \in (0, 1)$, and let $\{\mathbf{v}_1, \ldots, \mathbf{v}_M\}$ be a maximal δ-separated set for $\mathcal{U}$ with respect to $d_{n,\infty}(\cdot, \cdot)$. Note that $M = M_n^\infty(\mathcal{U}, \delta)$. Now let $\epsilon = (\delta/2)^{(1+p)/p}$, and let $\{\mathbf{u}_1, \ldots, \mathbf{u}_L\}$ be an ϵ-covering set for $\mathcal{U}$ with respect

to $d_{n,p}(\cdot,\cdot)$ with minimal cardinality. Note that $L = N(\mathcal{U},\epsilon,d_{n,p})$. By the union bound,

(3.4)

$$\begin{aligned} M_n^\infty(\mathcal{U},\delta) &= \left|\bigcup_{i=1}^{L} B_{n,p}(\mathbf{u}_i,\epsilon)\cap\{\mathbf{v}_1,\dots,\mathbf{v}_M\}\right| \\ &\le \sum_{i=1}^{L}\left|B_{n,p}(\mathbf{u}_i,\epsilon)\cap\{\mathbf{v}_1,\dots,\mathbf{v}_M\}\right| \\ &\le L\cdot\max\left\{\left|B_{n,p}(\mathbf{u}_i,\epsilon)\cap\{\mathbf{v}_1,\dots,\mathbf{v}_M\}\right| : i\in\{1,\dots,L\}\right\}. \end{aligned}$$

Applying Lemma 3.1 and the fact that $L = N(\mathcal{U},\epsilon,d_{n,p})$ in (3.4), we obtain

$$M_n^\infty(\mathcal{U},\delta) \le N(\mathcal{U},\epsilon,d_{n,p})(3K/\delta)^{\delta n/2}2^{H(\delta/2)n}. \tag{3.5}$$

Since any maximal δ-separated set must be a δ-covering set, we have $N(\mathcal{U},\delta,d_{n,\infty}) \le M_n^\infty(\mathcal{U},\delta)$, and then from (3.5), we see that

$$N(\mathcal{U},\delta,d_{n,\infty}) \le N(\mathcal{U},\epsilon,d_{n,p})(3K/\delta)^{\delta n/2}2^{H(\delta/2)n}.$$

Taking logarithm and dividing by n yields

$$\frac{1}{n}\log N(\mathcal{U},\delta,d_{n,\infty}) \le \frac{1}{n}\log N(\mathcal{U},\epsilon,d_{n,p}) + \frac{\delta}{2}\log(3K/\delta) + H(\delta/2)\log 2.$$

Thus letting n tend to infinity gives

$$h_\infty(\mathcal{U},\delta) \le h_p(\mathcal{U},\epsilon) + \frac{\delta}{2}\log(3K/\delta) + H(\delta/2)\log 2.$$

Since $\epsilon = (\delta/2)^{(1+p)/p}$, taking the limit as δ decreases to zero allows us to conclude that $h_\infty(\mathcal{D}) \le h_p(\mathcal{D})$. □

3.2. **Variational characterization of mean width.** Here we collect a few facts regarding mean width, which are used elsewhere. Let us begin with the fact that the sequence of finite sample mean widths is subadditive.

Remark 3.2. Using definition (1.3), one may easily check that that for $m,n \ge 1$,

$$\begin{aligned} \kappa_{m+n}(\mathcal{D}:\boldsymbol{\varepsilon}) &\le \mathbb{E}\left[\sup_{x,\theta}\sum_{k=0}^{m-1} f_\theta\circ T_\theta^k(x)\cdot\varepsilon_k\right] + \mathbb{E}\left[\sup_{x,\theta}\sum_{k=m}^{m+n-1} f_\theta\circ T_\theta^k(x)\cdot\varepsilon_k\right] \\ &\le \kappa_m(\mathcal{D}:\boldsymbol{\varepsilon}) + \kappa_n(\mathcal{D}:\boldsymbol{\varepsilon}). \end{aligned}$$

The last inequality above is a consequence of the stationarity of $\boldsymbol{\varepsilon}$ and the fact that $f_\theta\circ T_\theta^{k+m}(x) = f_\theta\circ T_\theta^k(T_\theta^m(x))$ for any $x\in\mathcal{X}$ and $\theta\in\Theta$.

Thus the sequence $\{\kappa_n(\mathcal{D} : \boldsymbol{\varepsilon}) : n \geq 1\}$ is subadditive, and therefore the limit in (1.4) exists.

The following result provides a variational characterization of the mean width.

Theorem 3.3. *If $\mathcal{D}$ is a family of dynamical models satisfying (D1)-(D3) and $\boldsymbol{\varepsilon}$ is a stationary ergodic process with finite mean, then*

$$\kappa(\mathcal{D} : \boldsymbol{\varepsilon}) \;=\; \sup_{\mathbf{U} \in \mathcal{Q}_{\mathcal{D}}} \sup_{\mathcal{J}(\mathbf{U}, \boldsymbol{\varepsilon})} \mathbb{E}\big[U_0 \cdot \varepsilon_0\big],$$

and the supremum is achieved.

Proof. Using the same system $(\mathcal{Z}, T)$ appearing in Section 2 as the model system, the noise process $\boldsymbol{\varepsilon}$ in place of the observation process, and a different cost function (product instead of loss), Theorem 3.3 is a consequence of [26, Theorem 1.4]. □

3.3. **Connecting entropy and mean width.** In this section we investigate connections between the notions of entropy and mean width for continuous families of dynamical models. We begin by proving that a family with zero entropy must have zero mean width relative to centered i.i.d. processes. Our proof relies on a foundational result of Furstenberg concerning joinings, stated below as Theorem B.

Let $\mathbf{U}$ be a stationary stochastic process taking values in a separable completely metrizable space $\mathcal{U}$. Let π be a finite Borel partition of $\mathcal{U}$. For $n \geq 1$, and $A_0^{n-1} = (A_0, \ldots, A_{n-1}) \in \pi^n$, let

$$p(A_0^{n-1}) = \mathbb{P}\big(U_0 \in A_0, \ldots, U_{n-1} \in A_{n-1}\big),$$

and consider

$$H_n(\mathbf{U}, \pi) = -\sum_{A_0^{n-1} \in \pi^n} p(A_0^{n-1}) \log p(A_0^{n-1}),$$

with the convention that $0 \cdot \log 0 = 0$. By subadditivity, we may take the limit as n tends to infinity:

$$h(\mathbf{U}, \pi) = \lim_n \frac{1}{n} H_n(\mathbf{U}, \pi).$$

Then taking the supremum over all finite Borel partitions of $\mathcal{U}$ gives the entropy of the process $\mathbf{U}$:

$$h(\mathbf{U}) = \sup_{\pi} h(\mathbf{U}, \pi).$$

In proving Theorem 1.24, we will rely on the following result of Furstenberg.

Theorem B. [8, Theorem I.2] *If $h(\mathbf{U}) = 0$ and $\mathbf{V}$ is i.i.d., then the only joining of $\mathbf{U}$ and $\mathbf{V}$ is the independent joining.*

Proof of Theorem 1.24. Let $\mathbf{U} \in \mathcal{Q}_{\mathcal{D}}$. Since $h(\mathcal{D}) = 0$, Lemma 1.17 implies that $h(\mathbf{U}) = 0$. Then by the result of Furstenberg (Theorem B above), the only joining of $\mathbf{U}$ with $\boldsymbol{\varepsilon}$ is the independent joining. Thus, for any joining $(\mathbf{U}, \boldsymbol{\varepsilon})$ with $\mathbf{U} \in \mathcal{Q}_{\mathcal{D}}$, we have

$$\mathbb{E}\big[U_0 \cdot \varepsilon_0\big] = \mathbb{E}\big[U_0\big] \cdot \mathbb{E}\big[\varepsilon_0\big] = 0.$$

Then by Theorem 3.3,

$$\kappa(\mathcal{D} : \boldsymbol{\varepsilon}) = \sup_{\mathbf{U} \in \mathcal{Q}_{\mathcal{D}}} \sup_{\mathcal{J}(\mathbf{U}, \boldsymbol{\varepsilon})} \mathbb{E}\big[U_0 \cdot \varepsilon_0\big] = 0.$$

Now suppose that $h(\mathcal{D}) > 0$ and that $(\varepsilon_k)_{k \geq 0}$ are i.i.d. standard normal random variables. For each $n \geq 1$, an application of Sudakov's lower bound [40] (see also [17, Theorem 6.1]) yields

$$\frac{1}{n} \kappa_n(\mathcal{D} : \boldsymbol{\varepsilon}) \geq \sup_{\delta > 0} \frac{\delta}{6} \left(\frac{1}{n} \log N(\mathcal{U}_{\mathcal{D}}, \delta, d_{n,2}) \right)^{1/2}.$$

By Theorem 1.13 and the positive entropy hypothesis, we have $h_2(\mathcal{U}_{\mathcal{D}}) = h(\mathcal{D}) > 0$. Therefore there exists $\delta > 0$ such that $h_2(\mathcal{U}_{\mathcal{D}}, \delta) > 0$. With this choice of δ, we take n to infinity in the previous display and obtain

$$\kappa_G(\mathcal{D}) \geq \frac{\delta}{6} \big(h_2(\mathcal{U}_{\mathcal{D}}, \delta) \big)^{1/2} > 0.$$

□

4. Signal plus noise

This section addresses empirical risk minimization for families dynamical models in the signal plus noise setting. Recall that joinings were defined for $\mathbb{R}$-valued processes in Definition 1.7.

The following result is an extension of Furstenberg's result, which we use to show that minimum risk estimates decouple the signal from the noise in the low complexity setting.

Theorem 4.1. *Suppose that $\mathbf{U}$ is a zero entropy process, $\mathbf{V}$ is a stationary ergodic process, and $\mathbf{W}$ is an i.i.d. process. Suppose that $(\mathbf{U}, \mathbf{V}, \mathbf{W})$ is a joining of these three processes such that $\mathbf{V}$ and $\mathbf{W}$ are independent. Then $(\mathbf{U}, \mathbf{V})$ is independent of $\mathbf{W}$.*

Note that Furstenberg's Theorem B is an immediate corollary of Theorem 4.1, obtained by taking $\mathbf{V}$ to be empty. To prove Theorem 4.1, we modify Furstenberg's proof of Theorem B. See Appendix A.2.

Our proofs also require the concept of the relatively independent joining, which results from a standard construction in ergodic theory [9, p. 126]. For notation, if $\mathbf{U}$ is a stationary processes taking values on $\mathcal{U}$ and $f : \mathcal{U} \to \mathcal{W}$ is a measurable map, then we let $f(\mathbf{U})$ denote the $\mathcal{W}$-valued process $\{f(U_k)\}_{k\geq 0}$, and we say that f maps $\mathbf{U}$ onto $\mathbf{W}$ whenever $f(\mathbf{U})$ has the same distribution as $\mathbf{W}$.

Proposition 4.2 (Relatively Indenpendent Joining)**.** *Suppose* $\mathbf{U}$*,* $\mathbf{V}$*, and* $\mathbf{W}$ *are stationary processes taking values on separable completely metrizable spaces. If there are Borel measurable maps* f *and* g *such that* $f(\mathbf{U})$ *and* $g(\mathbf{V})$ *each have the same distribution as* $\mathbf{W}$*, then there is a joining* $(\mathbf{U}, \mathbf{V})$ *of* $\mathbf{U}$ *and* $\mathbf{V}$ *such that* $f(\mathbf{U}) = g(\mathbf{V})$ *almost surely.*

The joining whose existence is asserted by Proposition 4.2 will be called the relatively independent joining of $\mathbf{U}$ and $\mathbf{V}$ (relative to $\mathbf{W}$). It will be used several times in the following proofs.

Proof of Theorem 1.18. By Theorem 1.12 any sequence of minimal ℓ-risk parameters converges almost surely to $\Theta_\ell(\mathbf{Y})$. We will complete the proof by showing that $\Theta_\ell(\mathbf{Y}) = \Theta_L(\mathbf{V})$.

Let $\mathbf{U}$ be in $\mathcal{Q}_\mathcal{D}$. Let $(\mathbf{U}, \mathbf{Y})$ be a joining of $\mathbf{U}$ and $\mathbf{Y}$ such that $\gamma_\ell(\mathbf{U}, \mathbf{Y}) = \mathbb{E}\big[\ell(U_0, Y_0)\big]$. Also, let $(\mathbf{V}, \boldsymbol{\varepsilon})$ be the independent joining of $\mathbf{V}$ and $\boldsymbol{\varepsilon}$. Using Proposition 4.2, let $(\mathbf{U}, \mathbf{Y}, \mathbf{V}, \boldsymbol{\varepsilon})$ be the relatively independent joining of $(\mathbf{U}, \mathbf{Y})$ and $(\mathbf{V}, \boldsymbol{\varepsilon})$ such that $\mathbf{Y} = \mathbf{V} + \boldsymbol{\varepsilon}$ almost surely. Projecting $(\mathbf{U}, \mathbf{Y}, \mathbf{V}, \boldsymbol{\varepsilon})$ onto the first, third, and fourth coordinates, we obtain a joining $(\mathbf{U}, \mathbf{V}, \boldsymbol{\varepsilon})$. By Theorem 4.1, the projection $(\mathbf{U}, \mathbf{V})$ is independent of $\boldsymbol{\varepsilon}$.

We have

$$\gamma_\ell(\mathbf{U}, \mathbf{Y}) = \mathbb{E}\big[\ell(U_0, Y_0)\big] = \mathbb{E}\big[\ell(U_0, V_0 + \varepsilon_0)\big].$$

Conditioning on $(\mathbf{U}, \mathbf{V})$ and using that these variables are jointly independent of $\boldsymbol{\varepsilon}$, we see that

$$\begin{aligned}\mathbb{E}\big[\ell(U_0, V_0 + \varepsilon_0)\big] &= \mathbb{E}\Big[\mathbb{E}\big[\ell(u, v + \varepsilon_0) \mid U_0 = u, V_0 = v\big]\Big] \\ &= \mathbb{E}\Big[L(U_0, V_0)\Big] \geq \gamma_L(\mathbf{U}, \mathbf{V}).\end{aligned}$$

Hence we have shown that $\gamma_\ell(\mathbf{U}, \mathbf{Y}) \geq \gamma_L(\mathbf{U}, \mathbf{V})$.

Now let $(\mathbf{U}, \mathbf{V})$ be a joining such that $\gamma_L(\mathbf{U}, \mathbf{V}) = \mathbb{E}\big[L(U_0, V_0)\big]$. Let $(\mathbf{U}, \mathbf{V}, \boldsymbol{\varepsilon})$ be the independent joining of $(\mathbf{U}, \mathbf{V})$ with $\boldsymbol{\varepsilon}$, and let

$\mathbf{Y} = \mathbf{V} + \boldsymbol{\varepsilon}$. Then

$$\begin{aligned}\gamma_L(\mathbf{U},\mathbf{V}) &= \mathbb{E}\Big[L(U_0,V_0)\Big] \\ &= \mathbb{E}\Big[\mathbb{E}\big[\ell(u,v+\varepsilon_0) \mid U_0=u, V_0=v\big]\Big] \\ &= \mathbb{E}\big[\ell(U_0,V_0+\varepsilon_0)\big] = \mathbb{E}\big[\ell(V_0,Y_0)\big] \geq \gamma_\ell(\mathbf{U},\mathbf{Y}).\end{aligned}$$

Thus $\gamma_L(\mathbf{U},\mathbf{V}) \geq \gamma_\ell(\mathbf{U},\mathbf{Y})$. Combining this inequality with the reverse inequality, which we established above, yields that $\gamma_L(\mathbf{U},\mathbf{V}) = \gamma_\ell(\mathbf{U},\mathbf{Y})$. As $\mathbf{U} \in \mathcal{Q}_\mathcal{D}$ was arbitrary, we obtain that $\Theta_\ell(\mathbf{Y}) = \Theta_L(\mathbf{V})$, which finishes the proof. □

Proof of Theorem 1.21. By Theorem 1.18 any sequence of minimum risk parameters converges almost surely to $\Theta_L(\mathbf{V})$. Assume for (1) that $\mathbb{E}\,\varepsilon_0 = 0$ and that $\ell(u,v) = D_F(v,u)$ is a Bregman divergence of a continuously differentiable convex function $F : \mathbb{R} \to \mathbb{R}$, namely

$$\ell(u,v) \;:=\; F(v) - F(u) - (v-u)F'(u).$$

Since $\mathbb{E}(\varepsilon) = 0$, we have

$$\begin{aligned}L(u,v) = \mathbb{E}\ell(u,v+\varepsilon_0) &= \mathbb{E}F(v+\varepsilon_0) - F(u) - (v-u)F'(u) \\ &= \ell(u,v) + G(v),\end{aligned}$$

where $G(v) = \mathbb{E}F(v+\varepsilon_0) - F(v)$ depends only on v and the distribution of ε_0 and is non-negative (since F is convex). Thus, for any $\theta \in \Theta$ and any $\mathbf{U} \in \mathcal{Q}_\theta$,

$$\begin{aligned}\gamma_L(\mathbf{U},\mathbf{V}) &= \inf_{\mathcal{J}(\mathbf{U},\mathbf{V})} \mathbb{E}L(U_0,V_0) \\ &= \inf_{\mathcal{J}(\mathbf{U},\mathbf{V})} \Big\{\mathbb{E}\ell(U_0,V_0) + \mathbb{E}G(V_0)\Big\} \\ &= \gamma_\ell(\mathbf{U},\mathbf{V}) + \mathbb{E}G(V_0).\end{aligned}$$

It follows that $\Theta_L(\mathbf{V}) = \Theta_\ell(\mathbf{V})$, as desired for (1).

Now assume for (2) that $\mathbf{V}$ is an ergodic process in $\mathcal{Q}_{\theta_0}$ and $\mathbb{E}\,\ell(x, y+\varepsilon_0) \geq \mathbb{E}\,\ell(0,\varepsilon_0)$ for all $x, y \in \mathbb{R}$, with equality if and only if $x = y$. Let $\Theta_1 = \{\theta \in \Theta : \mathbf{V} \in \mathcal{Q}_\theta\}$. We will finish the proof by showing that $\Theta_L(\mathbf{V}) = \Theta_1$.

Let $\mathbf{U} \in \mathcal{Q}_\mathcal{D}$, and let $(\mathbf{U},\mathbf{V})$ be a joining of these two processes. By the hypothesis concerning $\boldsymbol{\varepsilon}$ and ℓ, we have

$$\mathbb{E}\big[L(U_0,V_0)\big] = \mathbb{E}\Big[\mathbb{E}\big[\ell(u,v+\epsilon_0) \mid U_0=u, V_0=v\big]\Big] \tag{4.1}$$

$$\geq \mathbb{E}\big[\ell(0,\epsilon_0)\big], \tag{4.2}$$

with equality if and only if $U_0 = V_0$ almost surely. Since $(\mathbf{U}, \mathbf{V})$ is a joining (and in particular is stationary), we observe that $U_0 = V_0$ almost surely if and only if $\mathbf{U} = \mathbf{V}$ almost surely. Thus, we have shown that $\gamma_L(\mathbf{U}, \mathbf{V}) \geq \mathbb{E}\big[\ell(0, \epsilon_0)\big]$, with equality if and only if $\mathbf{U} = \mathbf{V}$. Therefore the set of θ minimizing the quantity $\min_{\mathbf{U} \in \mathcal{Q}_\theta} \gamma_L(\mathbf{U}, \mathbf{V})$ is exactly the set of θ such that $\mathbf{V} \in \mathcal{Q}_\theta$. Hence, $\Theta_L(\mathbf{V}) = \Theta_1$, which finishes the proof. □

4.1. **Least squares estimation.** Here we provide short proofs of our results pertaining to least squares estimation. It is possible to give somewhat more direct proofs of these results (avoiding Theorem 1.12, for example), but given the tools that we have already established, we present the most efficient proofs of which we are aware.

Proof of Theorem 1.26. By Theorem 1.12, any sequence of least squares parameters converges almost surely to $\Theta_\ell(\mathbf{Y})$. We will complete the proof by showing that $\Theta_\ell(\mathbf{Y}) = \operatorname{argmin}_\theta \min_{\mathbf{U} \in \mathcal{Q}_\theta} \gamma_2(\mathbf{U}, \mathbf{V})$.

Let $\theta \in \Theta$. Let $\mathbf{U}$ be in $\mathcal{Q}_\theta$, and let $(\mathbf{U}, \mathbf{Y})$ be a joining of these processes. Let $(\mathbf{V}, \boldsymbol{\varepsilon})$ be the independent joining, and (using Proposition 4.2) let $(\mathbf{U}, \mathbf{Y}, \mathbf{V}, \boldsymbol{\varepsilon})$ be a joining of $(\mathbf{U}, \mathbf{Y})$ with $(\mathbf{V}, \boldsymbol{\varepsilon})$ such that $\mathbf{Y} = \mathbf{U} + \boldsymbol{\varepsilon}$ almost surely.

Then

$$\begin{aligned} \mathbb{E}\Big[\big|U_0 - Y_0\big|^2\Big] &= \mathbb{E}\Big[\big|U_0 - (V_0 + \varepsilon_0)\big|^2\Big] \\ &= \mathbb{E}\Big[\big|U_0 - V_0\big|^2\Big] - 2\mathbb{E}\big[(U_0 - V_0) \cdot \varepsilon_0\big] + \mathbb{E}\Big[\varepsilon_0^2\Big]. \end{aligned} \tag{4.3}$$

Since $\kappa(\mathcal{D} : \boldsymbol{\varepsilon}) = 0$, we have $\mathbb{E}[U_0 \cdot \varepsilon_0] = 0$ by Theorem 3.3. Also, since V_0 is independent of ε_0 and ε_0 has zero mean, we have $\mathbb{E}[V_0 \cdot \varepsilon_0] = 0$. Applying these facts in (4.3), we see that

$$\mathbb{E}\Big[\big|U_0 - Y_0\big|^2\Big] = \mathbb{E}\Big[\big|U_0 - V_0\big|^2\Big] + \mathbb{E}\Big[\varepsilon_0^2\Big]. \tag{4.4}$$

Since $\mathbb{E}[\epsilon_0^2]$ is a constant that depends only on ϵ_0 (and not on the joining), we conclude that

$$\min_{\mathbf{U} \in \mathcal{Q}_\theta} \gamma_2(\mathbf{U}, \mathbf{Y}) \geq \min_{\mathbf{U} \in \mathcal{Q}_\theta} \gamma_2(\mathbf{U}, \mathbf{V}) + \mathbb{E}[\epsilon_0^2]. \tag{4.5}$$

Now let $\mathbf{U} \in \mathcal{Q}_\theta$ and $(\mathbf{U}, \mathbf{V})$ be a joining such that $\min_{\mathbf{U} \in \mathcal{Q}_\theta} \gamma_2(\mathbf{U}, \mathbf{V}) = \mathbb{E}[|U_0 - V_0|^2]$. Let $\boldsymbol{\varepsilon}$ be independent of $(\mathbf{U}, \mathbf{V})$, and let $\mathbf{Y} = \mathbf{V} + \boldsymbol{\varepsilon}$. Then $(\mathbf{U}, \mathbf{Y})$ is a joining of these processes such that

$$\mathbb{E}\Big[\big|U_0 - Y_0\big|^2\Big] = \mathbb{E}\Big[\big|U_0 - V_0\big|^2\Big] + \mathbb{E}\Big[\varepsilon_0^2\Big] = \min_{\mathbf{U} \in \mathcal{Q}_\theta} \gamma_2(\mathbf{U}, \mathbf{V}) + \mathbb{E}\Big[\varepsilon_0^2\Big]. \tag{4.6}$$

Combining (4.5) and (4.6), we obtain that $\min_{\mathbf{U}\in\mathcal{Q}_\theta}\gamma_2(\mathbf{U},\mathbf{Y})$ is equal to $\min_{\mathbf{U}\in\mathcal{Q}_\theta}\gamma_2(\mathbf{U},\mathbf{V})+\mathbb{E}\big[\varepsilon_0^2\big]$. Minimizing over θ yields that $\Theta_2(\mathbf{Y}) = \operatorname{argmin}_\theta \min_{\mathbf{U}\in\mathcal{Q}_\theta}\gamma_2(\mathbf{U},\mathbf{V})$, as was to be shown. □

Proof of Corollary 1.27. The hypothesis that each T_θ has zero entropy can be easily seen to be equivalent in this context to the statement that $h(\mathcal{D}) = 0$. Then by Theorem 1.21 any sequence of least squares parameters converges almost surely to the set $\{\theta\in\Theta : \mathbf{V}\in\mathcal{Q}_\theta\}$. Now suppose $\mathbf{V}\in\mathcal{Q}_\theta$. Then there exists a measure $\mu_0\in\mathcal{M}(\mathcal{X},T_\theta)$ such that if $\mathbf{U} = (T_\theta^k(X))_{k\geq 0}$, with X distributed according to μ_0, then $\mathbf{U}$ has the same distribution as $\mathbf{V}$. Hence X has the same distribution as V_0, which is given by μ, and therefore $\mu = \mu_0$. Furthermore, $(X, T_\theta(X))$ must have the same distribution as (V_0, V_1), which implies that $T_\theta(x) = T_{\theta^*}(x)$ for μ almost every x. We have thus shown that $\{\theta\in\Theta : \mathbf{V}\in\mathcal{Q}_\theta\}\subset\{\theta : \mu(T_\theta = T_{\theta^*}) = 0\}$. The reverse inclusion is obvious. □

5. Negative results

In this section, we show that least squares estimation can be almost surely *inconsistent* in a properly specified setting when $h(\mathcal{D}) > 0$ (or equivalently, $\kappa_G(\mathcal{D}) > 0$). We begin with a proof of Proposition 1.22.

Proof of Proposition 1.22. For the existence of a family $\mathcal{D}$ satisfying the hypotheses of the proposition, see Example 5.1. Suppose the hypotheses hold for a family $\mathcal{D}$ as in the statement of the proposition. In particular, fix an ergodic $\mathbf{V}\in\mathcal{Q}_{\theta_0}\setminus\cup_{\theta\neq\theta_0}\mathcal{Q}_\theta$.

By compactness and continuity, there exists K such that $|f_\theta(x)|\leq K$ for all $x\in\mathcal{X}$ and $\theta\in\Theta$. Since $h(\mathcal{D}) > 0$, Theorem 1.24 yields that $\kappa_G(\mathcal{D}) > 0$. Let $\sigma_0 = K^2/(2\kappa_G(\mathcal{D}))$. Let $\boldsymbol{\varepsilon}$ be an i.i.d. process of standard normal random variables, and let $\boldsymbol{\varepsilon}(\sigma) = \sigma\boldsymbol{\varepsilon}$. Now for $\sigma > \sigma_0$, let $\mathbf{Y} = \mathbf{V} + \boldsymbol{\varepsilon}(\sigma)$.

Let $\mathcal{D}_0 = \{(T_{\theta_0}, f_{\theta_0})\}$ be the family containing only the dynamical model corresponding to the parameter θ_0. By Lemma 1.17 and our hypothesis on θ_0, we have

$$h(\mathcal{D}_0) = \sup_{\mathbf{U}\in\mathcal{Q}_{\theta_0}} h(\mathbf{U}) = 0.$$

Then by (4.4), for any joining $(\mathbf{U},\mathbf{Y})$ with $\mathbf{U}\in\mathcal{Q}_{\theta_0}$,

$$\mathbb{E}\Big[\big|U_0 - Y_0\big|^2\Big] = \mathbb{E}\Big[\big|U_0 - V_0\big|^2\Big] + \mathbb{E}\Big[\varepsilon_0^2\Big] \geq \mathbb{E}\Big[\varepsilon_0^2\Big] = \sigma^2. \tag{5.1}$$

Hence, for any $\mathbf{U}\in\mathcal{Q}_{\theta_0}$, we have $\gamma_2(\mathbf{U},\mathbf{Y})\geq\sigma^2$.

Let us now show that there exists $\mathbf{U} \in \mathcal{Q}_{\mathcal{D}}$ such that $\gamma_2(\mathbf{U}, \mathbf{Y}) < \sigma^2$. As we have already established above, $\kappa_G(\mathcal{D}) > 0$. Then by Theorem 3.3, there exists a process $\mathbf{U} \in \mathcal{Q}_{\mathcal{D}}$ and a joining of $\mathbf{U}$ with $\boldsymbol{\varepsilon}$ such that

$$\mathbb{E}\Big[U_0 \cdot \varepsilon_0(\sigma)\Big] = \sigma \mathbb{E}\Big[U_0 \cdot \varepsilon_0\Big] = \sigma \kappa_G(\mathcal{D}).$$

Letting $(\mathbf{U}, \mathbf{V}, \boldsymbol{\varepsilon})$ be the independent joining of $\mathbf{V}$ with $(\mathbf{U}, \boldsymbol{\varepsilon})$, we have

$$\begin{aligned}
\gamma_2(\mathbf{U}, \mathbf{Y}) &\le \mathbb{E}\bigg[\Big|U_0 - (V_0 + \varepsilon_0)\Big|^2\bigg] \\
&= \mathbb{E}\Big[\big|U_0 - V_0\big|^2\Big] - 2\mathbb{E}\big[(U_0 - V_0) \cdot \varepsilon_0\big] + \mathbb{E}\Big[\varepsilon_0^2\Big] \\
&\le K^2 - 2\sigma\kappa_G(\mathcal{D}) + \sigma^2 < \sigma^2,
\end{aligned}$$

where the last inequality follows from our choice of $\sigma > \sigma_0$. Combining this inequality with (5.1), we conclude that θ_0 is not contained in $\Theta_2(\mathbf{Y})$, which is the nonempty compact limit set of least squares estimates by Theorem 1.12. □

The following example establishes the existence of a family $\mathcal{D}$ of dynamical models satisfying the hypotheses of Proposition 1.22. In fact, the parameter set Θ in this example only contains two values.

Example 5.1. Consider the state space $\mathcal{X} = [0, 1]$ with two transformations: let T_0 be the identity map on $[0, 1]$, and let T_1 be the fully chaotic logistic map, given by $T_1(x) = 4x(1 - x)$. We let $f_0 = f_1$ be the identity map on $[0, 1]$. This information fully determines $\mathcal{D}$.

As T_0 is the identity, we clearly have that $h(\mathbf{U}) = 0$ for all $\mathbf{U} \in \mathcal{Q}_{\theta_0}$. Also, we may let $\mathbf{V}$ be the constant process $V_k = 1/2$ for all k. Note that $\mathbf{V} \notin \mathcal{Q}_{\theta_1}$. Finally, it is well known that T_1 has positive entropy. We have thus verified the hypotheses of Proposition 1.22.

Appendix A. Additional proofs

A.1. **Proof of Lemma 1.17.** The well-known variational principle for topological dynamical systems [45, p. 190] equates the topological entropy $h_{\text{top}}(T)$ of a continuous transformation $T : \mathcal{X} \to \mathcal{X}$ of a compact metrizable space with the supremum of the measure-theoretic entropies $h_T(\mu)$ over $\mu \in \mathcal{M}(\mathcal{X}, T)$. Here we show how that variational principle may be applied in our setting to obtain Lemma 1.17. In the following lemma and its proof, we use subscripts to distinguish the different types of entropy under consideration. Here we write $h_{\text{fam}}(\mathcal{D})$ to denote the quantity $h(\mathcal{D})$ defined in Section 1.4, and we write $h_{\text{proc}}(\mathbf{U})$ to denote the quantity $h(\mathbf{U})$ defined in Section 3.3.

Lemma A.1. *For any continuous family $\mathcal{D}$ of dynamical models,*

$$h_{\text{fam}}(\mathcal{D}) = \sup_{\mathbf{U}\in\mathcal{Q}_{\mathcal{D}}} h_{\text{proc}}(\mathbf{U}).$$

Proof. Consider the dynamical system $(\mathcal{Z}, T)$ associated to $\mathcal{D}$ in Section 2, and let $\mathcal{U}$ be the set of sequences of the form $(f_\theta \circ T_\theta^k(x))_{k\geq 0}$ for any $\theta \in \Theta$ and $x \in \mathcal{X}$. By Lemma 2.1, the standard variational principle for topological dynamical systems may be applied, and therefore

$$h_{\text{top}}(T) = \sup_{\mu\in\mathcal{M}(\mathcal{Z},T)} h_T(\mu),$$

where the supremum may be taken over ergodic measures μ.

For θ in Θ, let $\mathcal{Z}_\theta = \text{proj}_\Theta^{-1}(\theta)$, which is closed and invariant under T. Also, let $T|_\theta$ denote the restriction of T to $\mathcal{Z}_\theta$.

For any scale $\delta > 0$, we have $N(\mathcal{U}, \delta, d_{n,\infty}) \leq N(\mathcal{Z}, \delta, d_{n,\infty})$, and therefore $h_{\text{fam}}(\mathcal{D}) \leq h_{\text{top}}(T)$. Similarly, we have that $h_{\text{top}}(T|_\theta) \leq h_{\text{fam}}(\mathcal{D})$.

By Lemma 2.1, for any ergodic $\mu \in \mathcal{M}(\mathcal{Z}, T)$, there exists θ such that $\mu = \delta_\theta \otimes \nu$, where ν is the distribution of a process $\mathbf{U} \in \mathcal{Q}_\theta$. Then $h_T(\mu) = h_{\text{proc}}(\mathbf{U})$, and we may conclude that

$$\sup_{\mu\in\mathcal{M}(\mathcal{Z},T)} h_T(\mu) \leq \sup_\theta \sup_{\mathbf{U}\in\mathcal{Q}_\theta} h_{\text{proc}}(\mathbf{U}).$$

Furthermore, for any $\theta \in \Theta$, an application of the standard variational principle to the system $(\mathcal{Z}_\theta, T|_\theta)$ gives that

$$h_{\text{top}}(T|_\theta) = \sup_{\mu\in\mathcal{M}(\mathcal{Z}_\theta,T|_\theta)} h_{T|_\theta}(\mu) = \sup_{\mathbf{U}\in\mathcal{Q}_\theta} h_{\text{proc}}(\mathbf{U}),$$

where the second equality follows from the equivalence between processes in $\mathcal{Q}_\theta$ and measures in $\mathcal{M}(\mathcal{Z}_\theta, T|_\theta)$. Combining all of the statements above, we see that

$$h_{\text{top}}(T) = \sup_{\mu\in\mathcal{M}(\mathcal{Z},T)} h_T(\mu) \leq \sup_{\mathbf{U}\in\mathcal{Q}_{\mathcal{D}}} h_{\text{proc}}(\mathbf{U}) \leq h_{\text{fam}}(\mathcal{D}) \leq h_{\text{top}}(T),$$

which concludes the proof of the lemma. □

A.2. **Proof of Theorem 4.1.** Here we provide the proof of Theorem 4.1, which is really an adaptation of Furstenberg's original proof of Theorem B. We begin with a lemma, which reduces the proof to the case of finite-valued processes.

Lemma A.2. *Suppose Theorem 4.1 holds for finite-valued processes. Then it holds in general.*

Proof. Suppose $\mathbf{U}$ takes values in the measurable space $(\mathcal{U}, \mathcal{F}_1)$, $\mathbf{V}$ in $(\mathcal{V}, \mathcal{F}_2)$, and $\mathbf{W}$ in $(\mathcal{W}, \mathcal{F}_3)$. Then the product σ-algebra on $\mathcal{U}^{\mathbb{N}} \times \mathcal{V}^{\mathbb{N}} \times \mathcal{W}^{\mathbb{N}}$ is generated by sets of the form $(A_1 \times \cdots \times A_N) \times (B_1 \times \cdots \times B_N) \times (C_1 \times \cdots \times C_N)$, where $A_i \in \mathcal{F}_1$, $B_i \in \mathcal{F}_2$, and $C_i \in \mathcal{F}_3$. Fix N and such $A_1, \ldots, A_N$, $B_1, \ldots, B_N$, and $C_1, \ldots, C_N$. Let $\pi_1 = \bigvee_i \{A_i, A_i^c\}$, $\pi_2 = \bigvee_i \{B_i, B_i^c\}$, and $\pi_3 = \bigvee_i \{C_i, C_i^c\}$, where $\bigvee$ denotes the join. Now consider the finite-valued processes $\pi_1(\mathbf{U})$, $\pi_2(\mathbf{V})$, and $\pi_3(\mathbf{W})$. Note that $\pi_1(\mathbf{U})$ has zero entropy, $\pi_2(\mathbf{V})$ is stationary and ergodic, and $\pi_3(\mathbf{W})$ is i.i.d. Furthermore, $\pi_2(\mathbf{V})$ is independent of $\pi_3(\mathbf{W})$. Hence, by Theorem 4.1 for finite-valued processes gives that the joint process $(\pi_1(\mathbf{U}), \pi_2(\mathbf{V}))$ is independent of $\pi_3(\mathbf{W})$. Then

$$\begin{aligned}
&\mathbb{P}\Big(\mathbf{U} \in A_1 \times \cdots \times A_N, \mathbf{V} \in B_1 \times \cdots \times B_N, \mathbf{W} \in C_1 \times \cdots \times C_N\Big) \\
&= \mathbb{P}\Big(\mathbf{U} \in A_1 \times \cdots \times A_N, \mathbf{V} \in B_1 \times \cdots \times B_N\Big)\mathbb{P}\Big(\mathbf{W} \in C_1 \times \cdots \times C_N\Big).
\end{aligned}$$

Since N and these sets were arbitrary, we see that $(\mathbf{U}, \mathbf{V})$ is independent of $\mathbf{W}$. □

Now, for the proof of Theorem 4.1, assume that all three processes are finite valued. In the rest of this proof, we use standard properties of the entropy of stationary ergodic processes (see [9]). First, notice that

$$\begin{aligned}
h(\mathbf{V}) + H(W_1) &= h((\mathbf{V}, \mathbf{W})) \\
&\leq h((\mathbf{U}, \mathbf{V}, \mathbf{W})) \\
&\leq h(\mathbf{U}) + h((\mathbf{V}, \mathbf{W})) \\
&= h(\mathbf{V}) + H(W_1),
\end{aligned}$$

where the two equalities rely our hypotheses: $h(\mathbf{U}) = 0$, $(\mathbf{V}, \mathbf{W})$ is the independent joining, and $\mathbf{W}$ is i.i.d. Then by conditioning and using

the information cocycle equation [9, p. 255], we have

$$\begin{aligned}
h(\mathbf{V}) + H(W_1) &\\
&= h((\mathbf{U}, \mathbf{V}, \mathbf{W})) \\
&= H(U_1, V_1, W_1 \mid U_2, V_2, W_2, \dots) \\
&= H(U_1, V_1 \mid U_2, V_2, W_2, \dots) + H(W_1 \mid U_1, V_1, U_2, V_2, W_2, \dots) \\
&\leq H(U_1 \mid U_2, V_2, W_2, \dots) + H(V_1 \mid U_2, V_2, W_2, \dots) \\
&\qquad\qquad + H(W_1 \mid U_1, V_1, U_2, V_2, W_2, \dots) \\
&\leq h(\mathbf{U}) + h(\mathbf{V}) + H(W_1 \mid U_1, V_1, U_2, V_2, W_2, \dots) \\
&= h(\mathbf{V}) + H(W_1 \mid U_1, V_1, U_2, V_2, W_2, \dots) \\
&\leq h(\mathbf{V}) + H(W_1).
\end{aligned}$$

Since $\mathbf{V}$ is finite-valued, $h(\mathbf{V})$ is finite. Then the previous display yields

$$H(W_1 \mid U_1, V_1, U_2, V_2, W_2, \dots) = H(W_1),$$

which implies that W_1 is independent of the σ-algebra generated by $\{U_1, V_1, U_2, V_2, \dots\}$. Repeating this argument for the variables $(W_1, \dots, W_r)$, $(W_{r+1}, \dots, W_{2r})$, $\dots$, shows that the entire process $\mathbf{W}$ is independent of $\{U_1, V_1, U_2, V_2, \dots\}$. This concludes the proof of the theorem.

References

[1] Terrence M Adams and Andrew B Nobel. Uniform convergence of Vapnik–Chervonenkis classes under ergodic sampling. *The Annals of Probability*, 38(4):1345–1367, 2010.

[2] Pierre Alquier and Olivier Wintenberger. Model selection for weakly dependent time series forecasting. *Bernoulli*, 18(3):883–913, 2012.

[3] Stéphane Boucheron, Gábor Lugosi, and Pascal Massart. *Concentration inequalities: A nonasymptotic theory of independence*. Oxford university press, 2013.

[4] Richard C Bradley. Basic properties of strong mixing conditions. a survey and some open questions. *Probability surveys*, 2(2):107–144, 2005.

[5] David Damanik, Mark Embree, and Anton Gorodetski. Spectral properties of schrödinger operators arising in the study of quasicrystals. In *Mathematics of Aperiodic Order*, pages 307–370. Springer, 2015.

[6] Thierry de la Rue. An introduction to joinings in ergodic theory. *Discrete and Continuous Dynamical Systems*, 15(1):121–142, 2006.

[7] Manfred Denker, Christian Grillenberger, and Karl Sigmund. *Ergodic theory on compact spaces*. Springer-Verlag, 1976.

[8] Harry Furstenberg. Disjointness in ergodic theory, minimal sets, and a problem in diophantine approximation. *Theory of Computing Systems*, 1(1):1–49, 1967.

[9] Eli Glasner. *Ergodic theory via joinings*. American Mathematical Soc., 2003.

[10] Robert M Gray, David L Neuhoff, and Paul C Shields. A generalization of Ornstein's d-bar distance with applications to information theory. *The Annals of Probability*, pages 315–328, 1975.

[11] H Hang and I Steinwart. A Bernstein-type inequality for some mixing processes and dynamical systems with an application to learning. *Ann. Statist.*, 2016.

[12] H Hang, I Steinwart, Y Feng, and J Suykens. Kernel density estimation for dynamical systems. *arXiv:1607.03792*, 2016.

[13] Michael P Hassell, John H Lawton, and RM May. Patterns of dynamical behaviour in single-species populations. *The Journal of Animal Ecology*, pages 471–486, 1976.

[14] Vitaly Kuznetsov and Mehryar Mohri. Time series prediction and online learning. In *29th Annual Conference on Learning Theory*, pages 1190–1213, 2016.

[15] Steven P. Lalley. Beneath the noise, chaos. *The Annals of Statistics*, 27(2):461–479, 1999.

[16] Steven P. Lalley and A. B. Nobel. Denoising deterministic time series. *Dyn. Partial Differ. Equ.*, 3(4):259–279, 2006.

[17] Michel Ledoux. Isoperimetry and gaussian analysis. In *Lectures on probability theory and statistics*, pages 165–294. Springer, 1996.

[18] Michel Ledoux and Michel Talagrand. *Probability in Banach Spaces: isoperimetry and processes*. Springer Science & Business Media, 2013.

[19] Nir Lev, Ron Peled, and Yuval Peres. Separating signal from noise. *Proceedings of the London Mathematical Society*, 110(4):883–931, 2015.

[20] Simon A Levin, Stephen R Carpenter, H Charles J Godfray, Ann P Kinzig, Michel Loreau, Jonathan B Losos, Brian Walker, and David S Wilcove. *The Princeton guide to ecology*. Princeton University Press, 2009.

[21] Daniel Marbach, James C Costello, Robert Küffner, Nicole M Vega, Robert J Prill, Diogo M Camacho, Kyle R Allison, Manolis Kellis, James J Collins, and Gustavo Stolovitzky. Wisdom of crowds for robust gene network inference. *Nature methods*, 9(8):796–804, 2012.

[22] Robert M May. Biological populations with nonoverlapping generations: stable points, stable cycles, and chaos. *Science*, 186(4164):645–647, 1974.

[23] K McGoff, X Guo, A Deckard, C Kelliher, A Leman, L Francey, J Hogenesch, S Haase, and J Harer. Local Edge Machine: inference of dynamic models of gene regulation. *Genome Biology*, 2016.

[24] Kevin McGoff, Sayan Mukherjee, Andrew Nobel, and Natesh Pillai. Consistency of maximum likelihood estimation for some dynamical systems. *The Annals of Statistics*, 43(1):1–29, 2015.

[25] Kevin McGoff, Sayan Mukherjee, and Natesh Pillai. Statistical inference for dynamical systems: A review. *Statist. Surv.*, 9:209–252, 2015.

[26] Kevin McGoff and Andrew Nobel. Optimal tracking for dynamical systems. arXiv:1601.05033, 2016.

[27] Gusztáv Morvai and Benjamin Weiss. On classifying processes. *Bernoulli*, pages 523–532, 2005.

[28] Gusztáv Morvai and Benjamin Weiss. Prediction for discrete time series. *Probability theory and related fields*, 132(1):1–12, 2005.

[29] Andrew B Nobel. Hypothesis testing for families of ergodic processes. *Bernoulli*, pages 251–269, 2006.

[30] Donald Ornstein. Bernoulli shifts with the same entropy are isomorphic. *Advances in Mathematics*, 4(3):337–352, 1970.

[31] Donald S Ornstein. An application of ergodic theory to probability theory. *The Annals of Probability*, 1(1):43–58, 1973.

[32] Donald S Ornstein. *Ergodic theory, randomness, and dynamical systems*. Yale University Press, 1974.

[33] Donald S Ornstein and Benjamin Weiss. How sampling reveals a process. *The Annals of Probability*, pages 905–930, 1990.

[34] David Pollard and Peter Radchenko. Nonlinear least-squares estimation. *Journal of Multivariate Analysis*, 97(2):548–562, 2006.

[35] E Arthur Robinson. The dynamical theory of tilings and quasicrystallography. *London Mathematical Society Lecture Note Series*, pages 451–474, 1996.

[36] Dan Shechtman, Ilan Blech, Denis Gratias, and John W Cahn. Metallic phase with long-range orientational order and no translational symmetry. *Physical Review Letters*, 53(20):1951, 1984.

[37] Paul Shields. *The ergodic theory of discrete sample paths*. American Mathematical Society, 1996.

[38] Boris Solomyak. Spectrum of dynamical systems arising from Delone sets. *Quasicrystals and Discrete Geometry, ed. J. Patera, Fields Institute Monographs*, 10:265–275, 1998.

[39] Ingo Steinwart and Marian Anghel. Consistency of support vector machines for forecasting the evolution of an unknown ergodic dynamical system from observations with unknown noise. *The Annals of Statistics*, pages 841–875, 2009.

[40] VN Sudakov. A remark on the criterion of continuity of gaussian sample function. In *Proceedings of the Second Japan-USSR Symposium on Probability Theory*, pages 444–454. Springer, 1973.

[41] Peter Turchin. *Complex population dynamics: a theoretical/empirical synthesis*, volume 35. Princeton University Press, 2013.

[42] Sara Van de Geer. Estimating a regression function. *The Annals of Statistics*, pages 907–924, 1990.

[43] V Vapnik and A Chervonenkis. On the uniform convergence of relative frequencies of events to their probabilities. *Theory of Probability & Its Applications*, 16(2):264–280, 1971.

[44] Divakar Viswanath, Xuan Liang, and Kirill Serkh. Metric entropy and the optimal prediction of chaotic signals. *SIAM Journal on Applied Dynamical Systems*, 12(2):1085–1113, 2013.

[45] Peter Walters. *An introduction to ergodic theory*, volume 79. Springer-Verlag, 1982.

[46] Chien-Fu Wu. Asymptotic theory of nonlinear least squares estimation. *The Annals of Statistics*, pages 501–513, 1981.

[47] A Zimin and C Lampert. Conditional risk minimization for stochastic processes. *arXiv:1510.02706*, 2016.